# DiBS: Differentiable Bayesian Structure Learning 

Lars Lorch<br>ETH Zurich<br>Zurich, Switzerland<br>lars.lorch@inf.ethz.ch

Jonas Rothfuss<br>ETH Zurich<br>Zurich, Switzerland<br>jonas.rothfuss@inf.ethz.ch

Bernhard Schölkopf<br>MPI for Intelligent Systems<br>Tübingen, Germany<br>bs@tuebingen.mpg.de

Andreas Krause<br>ETH Zurich<br>Zurich, Switzerland<br>krausea@ethz.ch


#### Abstract

Bayesian structure learning allows inferring Bayesian network structure from data while reasoning about the epistemic uncertainty-a key element towards enabling active causal discovery and designing interventions in real world systems. In this work, we propose a general, fully differentiable framework for Bayesian structure learning (DiBS) that operates in the continuous space of a latent probabilistic graph representation. Contrary to existing work, DiBS is agnostic to the form of the local conditional distributions and allows for joint posterior inference of both the graph structure and the conditional distribution parameters. This makes our formulation directly applicable to posterior inference of complex Bayesian network models, e.g., with nonlinear dependencies encoded by neural networks. Using DiBS, we devise an efficient, general purpose variational inference method for approximating distributions over structural models. In evaluations on simulated and real-world data, our method significantly outperforms related approaches to joint posterior inference. ${ }^{1}$


## 1 Introduction

Discovering the statistical and causal dependencies that underlie the variables of a data-generating system is of central scientific interest. Bayesian networks (BNs) [1] and structural equation models are commonly used for this purpose [2-5]. Structure learning, the task of learning a BN from observations of its variables, is well-studied, but computationally very challenging due to the combinatorially large number of candidate graphs and the constraint of graph acyclicity.

While structure learning methods arrive at a single plausible graph or its Markov equivalence class (MEC), e.g., [6-9], Bayesian structure learning aims to infer a full posterior distribution over BNs given the observations. A distribution over structures allows quantifying the epistemic uncertainty and the degree of confidence in any given BN model, e.g., when the amount of data is small. Most importantly, downstream tasks such as experimental design and active causal discovery rely on a posterior distribution over BNs to quantify the information gain from specific interventions and uncover the causal structure in a small number of experiments [10-15].
A key challenge in Bayesian structure learning is working with a posterior over BNs-a distribution over the joint space of (discrete) directed acyclic graphs and (continuous) conditional distribution parameters. Most of the practically viable approaches to Bayesian structure learning revolve around

[^0]Markov chain Monte Carlo (MCMC) sampling in combinatorial spaces and bootstrapping of classical score and constraint-based structure learning methods, e.g., in causal discovery [10-15]. However, these methods marginalize out the parameters and thus require a closed form for the marginal likelihood of the observations given the graph to remain tractable. This limits inference to simple and by now well-studied linear Gaussian and categorical BN models [16-18] and makes it difficult to infer more expressive BNs that, e.g., model nonlinear relationships among the variables. Due to the discrete nature of these approaches, recent advances in approximate inference and gradient-based optimization could not yet be translated into similar performance improvements in Bayesian structure learning.

In this work, we propose a novel, fully differentiable framework for Bayesian structure learning (DiBS) that operates in the continuous space of a latent probabilistic graph representation. Contrary to existing work, our formulation is agnostic to the distributional form of the BN and allows for inference of the joint posterior over both the conditional distribution parameters and the graph structure. This makes our approach directly applicable to more flexible BN models where neither the marginal likelihood nor the maximum likelihood parameter estimate have a closed form. We instantiate DiBS with the particle variational inference method of Liu and Wang [19] and present a general purpose method for approximate Bayesian structure learning. In our experiments on synthetic and real-world data, DiBS outperforms all alternative approaches to joint posterior inference of graphs and parameters and when modeling nonlinear interactions among the variables, often by a significant margin. This allows us to narrow down plausible causal graphs with greater precision and make better predictions under interventions-an important stepping stone towards active causal discovery

## 2 Background

Bayesian networks A Bayesian network $(\mathbf{G}, \boldsymbol{\Theta})$ models the joint density $p(\mathbf{x})$ of a set of $d$ variables $\mathbf{x}=x_{1: d}$ using (1) a directed acyclic graph (DAG) $\mathbf{G}$ encoding the conditional independencies of $\mathbf{x}$ and (2) parameters $\boldsymbol{\Theta}$ defining the local conditional distributions of each variable given its parents in the DAG. When modeling $p(\mathbf{x})$ using a BN, each variable is assumed to be independent of its non-descendants given its parents, thus allowing for a compact factorization of the joint $p(\mathbf{x} \mid \mathbf{G}, \boldsymbol{\Theta})$ into a product of local conditional distributions for each variable and its parents in $\mathbf{G}$.

Bayesian inference of BNs Given independent observations $\mathcal{D}=\left\{\mathbf{x}^{(1)}, \ldots, \mathbf{x}^{(N)}\right\}$, we consider the task of inferring a full posterior density over Bayesian networks that model the observations. Following Friedman and Koller [20], given a prior distribution over DAGs $p(\mathbf{G})$ and a prior over BN parameters $p(\boldsymbol{\Theta} \mid \mathbf{G})$, Bayes' Theorem yields the joint and marginal posterior distributions

$$
\begin{align*}
p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D}) & \propto p(\mathbf{G}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})  \tag{1}\\
p(\mathbf{G} \mid \mathcal{D}) & \propto p(\mathbf{G}) p(\mathcal{D} \mid \mathbf{G}) \tag{2}
\end{align*}
$$

where $p(\mathcal{D} \mid \mathbf{G})=\int p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) d \boldsymbol{\Theta}$ is the marginal likelihood. Thus, $p(\mathbf{G} \mid \mathcal{D})$ in (2) is only tractable in special conjugate cases where the integral over $\Theta$ can be computed in closed form. The Bayesian formalism allows us to compute expectations of the form

$$
\begin{equation*}
\mathbb{E}_{p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})}[f(\mathbf{G}, \boldsymbol{\Theta})] \quad \text { or } \quad \mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})}[f(\mathbf{G})] \tag{3}
\end{equation*}
$$

for any function $f$ of interest. For instance, to perform Bayesian model averaging, we would use $f(\mathbf{G}, \boldsymbol{\Theta})=p(\mathbf{x} \mid \mathbf{G}, \boldsymbol{\Theta})$ or $f(\mathbf{G})=p(\mathbf{x} \mid \mathbf{G})$, respectively [21, 22]. In active learning of causal BN structures, a commonly used $f$ is the expected decrease in entropy of $\mathbf{G}$ after an intervention [10-12, $14,15]$. Inferring either posterior is computationally challenging because there are $\mathcal{O}\left(d!2^{\binom{d}{2}}\right)$ possible DAGs with $d$ nodes [23]. Thus, computing the normalization constant $p(\mathcal{D})$ is generally intractable.

Continuous characterization of acyclic graphs Orthogonal to the work on Bayesian inference, Zheng et al. [9] have recently proposed a differentiable characterization of acyclic graphs for structure learning. In this work, we adopt the formulation of Yu et al. [24], who show that a graph with adjacency matrix $\mathbf{G} \in\{0,1\}^{d \times d}$ does not have any cycles if and only if $h(\mathbf{G})=0$, where

$$
\begin{equation*}
h(\mathbf{G}):=\operatorname{tr}\left[\left(\mathbf{I}+\frac{1}{d} \mathbf{G}\right)^{d}\right]-d . \tag{4}
\end{equation*}
$$

If $h(\mathbf{G})>0$, the function can be interpreted as quantifying the cyclicity or non-DAG-ness of $\mathbf{G}$. Follow-up work has leveraged this insight to model nonlinear relationships [24-27], time-series data [28], in the context of generative modeling [29, 30], for causal inference [31-33], and contributed to its theoretical understanding $[34,35]$. So far, a connection to Bayesian structure learning has been missing.

## 3 Related Work

The existing literature on Bayesian Structure Learning predominantly focuses on inferring the marginal graph posterior $p(\mathbf{G} \mid \mathcal{D})$. Since this requires $p(\mathcal{D} \mid \mathbf{G})$ to be tractable, inference is limited to BNs with linear Gaussian or Categorical conditional distributions [16-18]. By contrast, the formulation we introduce overcomes this fundamental restriction by allowing for joint inference of the graph and the parameters, thereby facilitating the active (causal) discovery of more expressive BNs.

MCMC Sampling from the posterior over graphs is the most general approach to approximate Bayesian structure learning. Structure MCMC $\left(\mathrm{MC}^{3}\right)$ [36, 37] performs Metropolis-Hastings in the space of DAGs by changing one edge at a time. Several works try to remedy its poor mixing behavior [38-40]. Alternatively, order MCMC draws samples in the smaller but still exponential space of node orders, which typically requires a hard limit on the maximum parent set size [20]. Attempts to correct for its unintended structural bias are themselves NP-hard to compute and/or limit the parent size [38, 41]. By performing variational inference in a continuous latent space, the method we propose circumvents such mixing issues and parent size limitations.

Bootstrapping The nonparametric DAG bootstrap [42] performs model averaging by bootstrapping $\mathcal{D}$, where each resampled data set is used to learn a single graph, e.g., using the GES or PC algorithms $[6,7]$. The obtained set of DAGs approximates the posterior by weighting each unique graph by its unnormalized posterior probability. In simple cases, a closed-form maximum likelihood parameter estimate may be used to approximate the joint posterior [14], but only if $p(\mathcal{D} \mid \mathbf{G})$ is tractable in the first place.

Exact methods A few notable exceptions use dynamic programming to achieve exact marginal inference in time $O\left(d 2^{d}\right)$, which is only feasible for $d \leq 20$ nodes [43, 44]. In special cases, e.g., for tree structures or known node orderings, exact inference can be performed more efficiently [45, 46].

## 4 A Fully Differentiable Framework for Bayesian Structure Learning

### 4.1 General Approach

With the goal of moving beyond the restrictive conjugate setups required by most discrete sampling methods for Bayesian structure learning, we propose to transfer the posterior inference task into the latent space of a probabilistic graph representation. Our resulting framework is consistent with the original Bayesian structure learning task in (3), enforces the acyclicity of $\mathbf{G}$ via the latent space prior, and provides the score of the continuous latent posterior, thus making general purpose inference methods applicable off-the-shelf.

Without loss of generality, we assume that there exists a latent variable $\mathbf{Z}$ that models the generative process of $\mathbf{G}$. Specifically, the default generative model described in Section 2 is generalized into the following factorization:


Figure 1: Generative model of BNs with latent variable $\mathbf{Z}$. This formulation generalizes the standard Bayesian setup in (1) where only $\mathbf{G}, \boldsymbol{\Theta}$, and x are modeled explicitly.

$$
\begin{equation*}
p(\mathbf{Z}, \mathbf{G}, \boldsymbol{\Theta}, \mathcal{D})=p(\mathbf{Z}) p(\mathbf{G} \mid \mathbf{Z}) p(\mathbf{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) \tag{5}
\end{equation*}
$$

Figure 1 displays the corresponding graphical model. The following insight provides us with an equivalence between the expectation we ultimately want to approximate and an expectation over the posterior of the continuous variable $\mathbf{Z}$ :
Proposition 1 (Latent posterior expectation). Under the generative model in (5), it holds that
(a) $\quad \mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})}[f(\mathbf{G})]=\mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}) p(\mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]}\right]$, and
(b) $\mathbb{E}_{p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})}[f(\mathbf{G}, \boldsymbol{\Theta})]=\mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}, \boldsymbol{\Theta}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}\right]$.

A proof is provided in Appendix A.1. Rather than approximating $p(\mathbf{G} \mid \mathcal{D})$ or $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$, our goal will be to infer $p(\mathbf{Z} \mid \mathcal{D})$ or $p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})$ instead, which by Proposition 1 allows us to compute
expectations of the form in (3). In the following, we first discuss how to define the two factors $p(\mathbf{G} \mid \mathbf{Z})$ and $p(\mathbf{Z})$ in a way that models only directed acyclic graphs. Then, we provide the details necessary to perform black box inference of the posteriors of the continuous latent variable $\mathbf{Z}$.

### 4.2 Representing DAGs in a Continuous Latent Space

Generative model of directed graphs We define the latent variable $\mathbf{Z}$ as consisting of two embedding matrices $\mathbf{U}, \mathbf{V} \in \mathbb{R}^{k \times d}$, i.e., $\mathbf{Z}=[\mathbf{U}, \mathbf{V}]$. Building on Kipf and Welling [47], we propose to use a bilinear generative model for the adjacency matrix $\mathbf{G} \in\{0,1\}^{d \times d}$ of directed graphs using the inner product between the latent variables in $\mathbf{Z}$ :

$$
\begin{equation*}
p_{\alpha}(\mathbf{G} \mid \mathbf{Z})=\prod_{i=1}^{d} \prod_{j \neq i}^{d} p_{\alpha}\left(g_{i j} \mid \mathbf{u}_{i}, \mathbf{v}_{j}\right) \text { with } p_{\alpha}\left(g_{i j}=1 \mid \mathbf{u}_{i}, \mathbf{v}_{j}\right)=\sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right) \tag{6}
\end{equation*}
$$

where $\sigma_{\alpha}(x)=1 /(1+\exp (-\alpha x))$ denotes the sigmoid function with inverse temperature $\alpha$. We denote the corresponding matrix of edge probabilities in $\mathbf{G}$ given $\mathbf{Z}$ by $\mathbf{G}_{\alpha}(\mathbf{Z}) \in[0,1]^{d \times d}$ with

$$
\begin{equation*}
\mathbf{G}_{\alpha}(\mathbf{Z})_{i j}:=p_{\alpha}\left(g_{i j}=1 \mid \mathbf{u}_{i}, \mathbf{v}_{j}\right) . \tag{7}
\end{equation*}
$$

Since we model acyclic graphs, which do not contain self-loops, we set $p_{\alpha}\left(g_{i i}=1 \mid \mathbf{Z}\right):=0$. The latent dimensionality $k$ trades off the complexity of the variable interactions with tractability during inference. For $k \geq d$, the matrix of edge probabilities $\mathbf{G}_{\alpha}(\mathbf{Z})$ is not constrained in rank, and the generative model in (6) can represent any adjacency matrix without self-loops. That said, the size of the latent representation $\mathbf{Z}$ only grows $\mathcal{O}(d \cdot k)$ and, in principle, $k$ can be chosen independently of $d$. Note that the formulation in (6) models directed graphs since $\sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right) \neq \sigma_{\alpha}\left(\mathbf{u}_{j}^{\top} \mathbf{v}_{i}\right)$. We can even interpret $\mathbf{u}_{i}$ and $\mathbf{v}_{i}$ as node embeddings that may encode more information than mere edge probabilities, e.g., for graph neural networks. The fact that $p_{\alpha}(\mathbf{G} \mid \mathbf{Z})$ is invariant to orthogonal transformations of $\mathbf{U}$ and $\mathbf{V}$ is not an issue in practice.
Acyclicity via the latent prior distribution A major constraint in learning BNs is the acyclicity of $\mathbf{G}$. With a latent graph model $p(\mathbf{G} \mid \mathbf{Z})$ in place, we design the prior $p(\mathbf{Z})$ to act as a soft constraint enforcing that only DAGs are modeled. Specifically, we define the prior of $\mathbf{Z}$ as the product of independent Gaussians with a Gibbs distribution that penalizes the expected cyclicity of $\mathbf{G}$ given $\mathbf{Z}$ :

$$
\begin{equation*}
p_{\beta}(\mathbf{Z}) \propto \exp \left(-\beta \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[h(\mathbf{G})]\right) \prod_{i j} \mathcal{N}\left(z_{i j} ; 0, \sigma_{z}^{2}\right) \tag{8}
\end{equation*}
$$

Here, $h$ is the DAG constraint function given in (4) and $\beta$ is an inverse temperature parameter controlling how strongly the acyclicity is enforced. As $\beta \rightarrow \infty$, the support of $p_{\beta}(\mathbf{Z})$ reduces to all $\mathbf{Z}$ that only model valid DAGs. The Gaussian component ensures that the norm of $\mathbf{Z}$ is well-behaved. Traditional graph priors of the form $p(\mathbf{G}) \propto f(\mathbf{G})$ that induce, e.g., sparsity of $\mathbf{G}$, can be flexibly incorporated into $p_{\beta}(\mathbf{Z})$ by means of an additional factor involving $f\left(\mathbf{G}_{\alpha}(\mathbf{Z})\right)$ or $\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})]$. Unless highlighting a specific point, we omit writing the hyperparameters $\alpha$ and $\beta$ to simplify notation.

### 4.3 Estimators for Gradient-Based Bayesian Inference

The extended generative model in (5) not only allows us to incorporate the notoriously difficult acyclicity constraint into the Bayesian framework. By rephrasing the posteriors $p(\mathbf{G} \mid \mathcal{D})$ and $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$ in terms of $p(\mathbf{Z} \mid D)$ and $p(\mathbf{Z}, \boldsymbol{\Theta} \mid D)$, respectively, it also makes Bayesian structure learning amenable to variational inference techniques that operate in continuous space and rely on the gradient of the unnormalized log posterior, also known as the score. The following result provides us with the score of both latent posteriors. Detailed derivations can be found in Appendix A.2.
Proposition 2 (Latent posterior score). Under the generative graph model defined in (5), the gradient of the log posterior density of $\mathbf{Z}$ is given by

$$
\begin{equation*}
\text { (a) } \quad \nabla_{\mathbf{Z}} \log p(\mathbf{Z} \mid \mathcal{D})=\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]} \tag{9}
\end{equation*}
$$

which is relevant when the marginal likelihood $p(\mathcal{D} \mid \mathbf{G})$ can be computed efficiently. In the general case, when inferring the joint posterior of $\mathbf{Z}$ and $\mathbf{\Theta}$, the gradients of the $\log$ posterior are given by

$$
\begin{equation*}
\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})=\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})]} \tag{b}
\end{equation*}
$$

(c) $\quad \nabla_{\boldsymbol{\Theta}} \log p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})=\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}\left[\nabla_{\boldsymbol{\Theta}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})\right]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})]}$
where $p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})=p(\mathbf{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \mathbf{\Theta})$, which is efficient to compute by construction.
The expectations have tractable Monte Carlo approximations because sampling from $p(\mathbf{G} \mid \mathbf{Z})$ in (6) is simple and parallelizable. The gradient terms of the form $\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[\cdot]$, which also appear inside $\nabla_{\mathbf{Z}} \log p(\mathbf{Z})$, can be estimated using two different techniques, depending on the BN model inferred.
Differentiable (marginal) likelihood Using the Gumbel-softmax trick [48, 49], we can separate the randomness from $\mathbf{Z}$ when sampling from $p(\mathbf{G} \mid \mathbf{Z})$ and obtain the following estimator:

$$
\begin{equation*}
\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})] \approx \mathbb{E}_{p(\mathbf{L})}\left[\left.\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})\right|_{\mathbf{G}=\mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})} \cdot \nabla_{\mathbf{Z}} \mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})\right] \tag{12}
\end{equation*}
$$

where $\mathbf{L} \sim \operatorname{Logistic}(0,1)^{d \times d}$ i.i.d.. The matrix-valued function $\mathbf{G}_{\tau}(\cdot)$ is defined elementwise as

$$
\mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})_{i j}:= \begin{cases}\sigma_{\tau}\left(l_{i j}+\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right) & \text { if } i \neq j  \tag{13}\\ 0 & \text { if } i=j\end{cases}
$$

The estimator applies equally to $p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ in place of $p(\mathcal{D} \mid \mathbf{G})$. For the estimator to be welldefined, $p(\mathcal{D} \mid \mathbf{G})$ or $p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$, respectively, needs to be differentiable with respect to $\mathbf{G}$. More specifically, the gradient $\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})$ or $\nabla_{\mathbf{G}} p(\mathcal{D}, \boldsymbol{\Theta} \mid \mathbf{G})$ needs to be defined when $\mathbf{G}$ lies on the interior of $[0,1]^{d \times d}$ and not at its discrete endpoints. This depends on the parameterization of the BN model we want to infer. In case $p(\mathcal{D} \mid \mathbf{G})$ or $p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ is only defined for discrete $\mathbf{G}$, it is possible to evaluate $\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})$ or $\nabla_{\mathbf{G}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ using hard Gumbel-max samples of $\mathbf{G}$ (i.e., with $\tau=\infty$ ) and a straight-through gradient estimator. Since the DAG constraint $h$ is differentiable, the Gumbel-softmax trick can always be applied inside $\nabla_{\mathbf{Z}} \log p(\mathbf{Z})$. In practice, we always use $\tau=1$.

Non-differentiable (marginal) likelihood In general, $\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})$ or $\nabla_{\mathbf{G}} p(\mathcal{D}, \boldsymbol{\Theta} \mid \mathbf{G})$ depending on the inference task might be not available or ill-defined. In this setting, the score function estimator provides us with a way to estimate the gradient we need [50]:

$$
\begin{equation*}
\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]=\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}\left[(p(\mathcal{D} \mid \mathbf{G})-b) \nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})\right] \tag{14}
\end{equation*}
$$

The estimator likewise applies for $p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})$ in place of $p(\mathcal{D} \mid \mathbf{G})$. Here, $b$ is a constant with respect to $\mathbf{G}$ that can be used for variance reduction [51], and $\nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})$ is trivial to compute. The derivations of both (12) and (14), alongside a more detailed discussion, can be found in Appendix B.

## 5 Particle Variational Inference for Structure Learning

In the previous section, we have proposed a differentiable formulation for Bayesian structure learning ( DiBS ) that is agnostic to the form of the local BN conditionals and, more importantly, translates learning discrete graph structures into an inference problem over the continuous variable $\mathbf{Z}$. In the following, we overcome the remaining challenge of inferring the intractable DiBS posteriors $p(\mathbf{Z} \mid \mathcal{D})$ and $p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})$ by employing Stein variational gradient descent (SVGD) [19], a gradient-based and general purpose variational inference method. The resulting algorithm infers a particle approximation of the marginal or joint posterior density over BNs given observational data.
SVGD for posterior inference Since Proposition 2 provides us with the gradient of the latent posterior score functions, we can apply SVGD off-the-shelf. SVGD minimizes the KL divergence to a target distribution by iteratively transporting a set of particles using a sequence of kernel-based transformation steps. We provide a more detailed overview of SVGD in Appendix C. Following this paradigm for DiBS, we iteratively update a fixed set $\left\{\mathbf{Z}^{(m)}\right\}_{m=1}^{M}$ or $\left\{\mathbf{Z}^{(m)}, \boldsymbol{\Theta}^{(m)}\right\}_{m=1}^{M}$ to approximate $p(\mathbf{Z} \mid \mathcal{D})$ or $p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})$, respectively. If the BN model we aim to infer has a properly-defined likelihood gradient with respect to $\mathbf{G}$, we use the Gumbel-softmax estimator in (12) to approximate the posterior score. Otherwise, we resort to the score function estimator in (14). We use a simple kernel for SVGD:

$$
\begin{equation*}
k\left((\mathbf{Z}, \boldsymbol{\Theta}),\left(\mathbf{Z}^{\prime}, \boldsymbol{\Theta}^{\prime}\right)\right):=\exp \left(-\frac{1}{\gamma_{z}}\left\|\mathbf{Z}-\mathbf{Z}^{\prime}\right\|_{2}^{2}\right)+\exp \left(-\frac{1}{\gamma_{\theta}}\left\|\boldsymbol{\Theta}-\boldsymbol{\Theta}^{\prime}\right\|_{2}^{2}\right) \tag{15}
\end{equation*}
$$

with bandwidths $\gamma_{z}, \gamma_{\theta}$. For inference of $p(\mathbf{Z} \mid \mathcal{D})$, we leave out the second term involving $\boldsymbol{\Theta}, \boldsymbol{\Theta}^{\prime}$. While $\mathbf{Z}$ is invariant to orthogonal transformations, more elaborate kernels that are, e.g., invariant to such transforms empirically perform worse in our experiments.

```
Algorithm 1 DiBS with SVGD [19] for inference of \(p(\mathbf{G} \mid \mathcal{D})\)
    Input: Initial set of latent particles \(\left\{\mathbf{Z}_{0}^{(m)}\right\}_{m=1}^{M}\), kernel \(k\), schedules for \(\alpha_{t}, \beta_{t}\), and stepsizes \(\eta_{t}\)
    Output: Set of discrete graph particles \(\left\{\mathbf{G}^{(m)}\right\}_{m=1}^{M}\) approximating \(p(\mathbf{G} \mid \mathcal{D})\)
    Incorporate prior belief of \(p(\mathbf{G})\) into \(p(\mathbf{Z}) \quad \triangleright\) See Section 4.2
    for iteration \(t=0\) to \(T-1\) do
        Estimate score \(\nabla_{\mathbf{Z}} \log p(\mathbf{Z} \mid \mathcal{D})\) given in (9) for each \(\mathbf{Z}_{t}^{(m)} \quad \triangleright\) See (12) and (14)
        for particle \(m=1\) to \(M\) do
            \(\mathbf{Z}_{t+1}^{(m)} \leftarrow \mathbf{Z}_{t}^{(m)}+\eta_{t} \phi_{t}\left(\mathbf{Z}_{t}^{(m)}\right) \quad \triangleright\) SVGD step
            where \(\boldsymbol{\phi}_{t}(\cdot):=\frac{1}{M} \sum_{k=1}^{M}\left[k\left(\mathbf{Z}_{t}^{(k)}, \cdot\right) \nabla_{\mathbf{Z}_{t}^{(k)}} \log p\left(\mathbf{Z}_{t}^{(k)} \mid \mathcal{D}\right)+\nabla_{\mathbf{z}_{t}^{(k)}} k\left(\mathbf{Z}_{t}^{(k)}, \cdot\right)\right]\)
    return \(\left\{\mathbf{G}_{\infty}\left(\mathbf{Z}_{T}^{(m)}\right)\right\}_{m=1}^{M} \quad \triangleright\) See (16) and (17)
```

Annealing $\alpha$ and $\beta$ The latent variable $\mathbf{Z}$ not only probabilistically models the graph $\mathbf{G}$, but can also be viewed as a continuous relaxation of $\mathbf{G}$, with $\alpha$ trading off smoothness with accuracy. As $\alpha \rightarrow \infty$, the sigmoid $\sigma_{\alpha}(\cdot)$ converges to the unit step function. Hence, as $\alpha \rightarrow \infty$ in the graph model $p_{\alpha}(\mathbf{G} \mid \mathbf{Z})$ in (6), the expectations in Proposition 1 simplify to:

$$
\begin{align*}
\mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})}[f(\mathbf{G})] & \rightarrow \mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[f\left(\mathbf{G}_{\infty}(\mathbf{Z})\right)\right] \\
\mathbb{E}_{p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})}[f(\mathbf{G}, \boldsymbol{\Theta})] & \rightarrow \mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[f\left(\mathbf{G}_{\infty}(\mathbf{Z}), \boldsymbol{\Theta}\right)\right] \tag{16}
\end{align*}
$$

where $\mathbf{G}_{\infty}(\mathbf{Z})$ denotes the single limiting graph implied by $\mathbf{Z}=[\mathbf{U}, \mathbf{V}]$ and is defined as

$$
\mathbf{G}_{\infty}(\mathbf{Z})_{i j}:= \begin{cases}1 & \text { if } \mathbf{u}_{i}^{\top} \mathbf{v}_{j}>0 \text { and } i \neq j  \tag{17}\\ 0 & \text { otherwise }\end{cases}
$$

See Appendix A.3. In this case, $p_{\alpha}(\mathbf{G} \mid \mathbf{Z})$ converges to representing only a single graph. To be able to invoke this simplification, we anneal $\alpha \rightarrow \infty$ over the iterations of SVGD and, upon termination, convert the latent variables $\mathbf{Z}$ to the single discrete $\mathbf{G}_{\infty}(\mathbf{Z})$. Furthermore, we similarly let $\beta \rightarrow \infty$ in the latent prior $p_{\beta}(\mathbf{Z})$ over the iterations to enforce that the latent representation of $\mathbf{G}$ only models DAGs. By Equation (16), the resulting DAGs form a consistent particle approximation of $p(\mathbf{G} \mid \mathcal{D})$ or $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$, respectively. Algorithm 1 summarizes DiBS instantiated with SVGD for inference of $p(\mathbf{G} \mid \mathcal{D})$. The general case of inferring $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$ is given in Algorithm 2 of Appendix D .

Single-particle approximation SVGD reduces to regular gradient ascent for the maximum a posteriori estimate when transporting only a single particle [19]. In this special case, DiBS with SVGD recovers some of the existing continuous structure learning methods: gradient ascent on a linear Gaussian likelihood solves an optimization problem similar to NOTEARS [9]. The cyclicity penalizer acts analogously. However, not only does DiBS automatically turn into a full Bayesian approach when using more particles, it is also not limited to settings such as linear Gaussian conditionals, where the adjacency matrix $\mathbf{G}$ and the parameters $\boldsymbol{\Theta}$ can be modeled together by a weighted adjacency matrix.

Weighted particle mixture In high dimensional settings, it may be beneficial to move beyond a uniform weighting of the inferred particles of BN models to approximate $p(\mathbf{G} \mid \mathcal{D})$ or $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$. We consider a particle mixture that weights each particle by its unnormalized posterior probability $p(\mathbf{G}, \mathcal{D})$ or $p(\mathbf{G}, \boldsymbol{\Theta}, \mathcal{D})$, respectively, under the BN model. While we do not have a strong theoretical justification for the weighting, our motivation is that most of the particles in the empirical distribution will be unique as a consequence of the super-exponentially large space of DAGs, which may result in a crude approximation of their posterior probability mass function. In our experiments, DiBS and its instantiation with SVGD are used interchangeably, and DiBS+ denotes the weighted particle mixture.

## 6 Evaluation on Synthetic Data

### 6.1 Experimental Setup

Synthetic data We compare DiBS to a set of related methods in marginal and joint posterior inference of synthetic linear and nonlinear Gaussian BNs. Our setup follows [9, 24, 27, 52], who consider inferring BNs with Erdős-Rényi and scale-free random structures [53, 54]. For each graph,


Figure 2: Marginal posterior inference of 20-node linear Gaussian BNs using the BGe marginal likelihood. Higher scores on AUROC and lower scores on $\mathbb{E}-\mathrm{SHD}$, neg. MLL, neg. I-MLL are preferred. DiBS+ performs competitively across all metrics, in particular $\mathbb{E}$-SHD and neg. (I-)MLL.
here with $d \in\{20,50\}$ nodes and $2 d$ edges in expectation, we sample a set of ground truth parameters and then generate training, held-out, and interventional data sets. In all settings, we use $N=100$ observations for inference, emulating the use case of Bayesian structure learning where the uncertainty about the graph structure is significant.
Graph priors For Erdős-Rényi graphs, all methods use the prior $p(\mathbf{G}) \propto q^{\|\mathbf{G}\|_{1}}(1-q)^{\binom{d}{2}-\|\mathbf{G}\|_{1}}$, capturing that each edge exists independently w.p. $q$ [53]. For scale-free graphs, we define the prior $p(\mathbf{G}) \propto \prod_{i=1}^{d}\left(1+\left\|\mathbf{G}_{i}^{\top}\right\|_{1}\right)^{-3}$, analogous to their power law degree distribution $p(\mathrm{deg}) \sim \operatorname{deg}^{-3}$ [54]. Here, $\mathbf{G}_{i}^{\top}$ is the $i$-th column of the adjacency matrix. DiBS implements either prior by using the corresponding term above as an additional factor in $p(\mathbf{Z})$ with $\mathbf{G}:=\mathbf{G}_{\alpha}(\mathbf{Z})$ (see Section 4.2).
Metrics Since neither density nor samples of the ground truth posteriors are available for BNs of $d \in\{20,50\}$ variables, we follow the evaluation metrics used by previous work. We define the expected structural Hamming distance to the ground truth graph $\mathbf{G}^{*}$ under the inferred posterior as

$$
\begin{equation*}
\mathbb{E}-\operatorname{SHD}\left(p, \mathbf{G}^{*}\right):=\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathcal{D}) \cdot \operatorname{SHD}\left(\mathbf{G}, \mathbf{G}^{*}\right) \tag{18}
\end{equation*}
$$

where $\operatorname{SHD}\left(\mathbf{G}, \mathbf{G}^{*}\right)$ counts the edge changes that separate the essential graphs representing the MECs of $\mathbf{G}$ and $\mathbf{G}^{*}$ [8, 55]. In addition, we follow Friedman and Koller [20] and Ellis and Wong [41] and compute the area under the receiver operating characteristic curve (AUROC) for pairwise edge predictions when varying the confidence threshold under the inferred marginal $p\left(g_{i j}=1 \mid \mathcal{D}\right)$. Finally, following Murphy [10], we also evaluate the ability to predict future observations by computing the average negative (marginal) log likelihood on 100 held-out observations $\mathcal{D}^{\text {test }}$ :

$$
\begin{equation*}
\text { neg. } \operatorname{LL}\left(p, \mathcal{D}^{\text {test }}\right):=-\sum_{\mathbf{G}, \boldsymbol{\Theta}} p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D}) \cdot \log p\left(\mathcal{D}^{\text {test }} \mid \mathbf{G}, \boldsymbol{\Theta}\right) \tag{19}
\end{equation*}
$$

When inferring $p(\mathbf{G} \mid \mathcal{D})$, the corresponding neg. MLL metric instead uses $p\left(\mathcal{D}^{\text {test }} \mid \mathbf{G}\right)$. Analogously, we also compute the interventional log likelihoods I-LL and I-MLL, a relevant metric in causal inference $[10,12]$. Here, an interventional data set $\left(\mathcal{D}^{\text {int }}, \mathcal{I}\right)$ is instead used to compute $p\left(\mathcal{D}^{\text {int }} \mid \mathbf{G}, \boldsymbol{\Theta}, \mathcal{I}\right)$ and $p\left(\mathcal{D}^{\text {int }} \mid \mathbf{G}, \mathcal{I}\right)$ in (19), respectively. Scores are the average of 10 interventional data sets. All reported metrics in this section are aggregated for inference of 30 random synthetic BNs.
In the remainder, $\operatorname{DiBS}$ is always run for 3,000 iterations and with $k=d$ for inference of $d$-variable BNs, which leaves the matrix of edge probabilities unconstrained in rank. We discard a DiBS particle in the rare case that a returned graph is cyclic. Complete details on Gaussian BNs, the evaluation metrics, hyperparameters, and all baselines can be found in Appendix E.

### 6.2 Linear Gaussian Bayesian Networks

Marginal posterior inference For linear Gaussian BNs, we first evaluate the classical setting of inferring the marginal posterior $p(\mathbf{G} \mid \mathcal{D})$ since $p(\mathcal{D} \mid \mathbf{G})$ can be computed in closed form. To this end, we employ the commonly used Bayesian Gaussian Equivalent (BGe) marginal likelihood, which


Figure 3: Joint posterior inference of graphs and parameters of linear Gaussian networks with $d=20$ nodes. DiBS+ performs best across all of the metrics. BGES*, the next-best alternative, yields substantially worse performance in $\mathbb{E}$-SHD, i.e., in recovering the overall graph structure and MEC. Recall that higher AUROC and lower $\mathbb{E}-\mathrm{SHD}$, neg. LL, and neg. I-LL scores are preferable.
scores Markov equivalent structures equally [16, 17]. The form of the BGe score requires DiBS to use the score function estimator in (14).
We compare DiBS with the nonparametric DAG bootstrap [42] using the constraint-based PC [7] and the score-based GES [6] algorithms (BPC and BGES). For MCMC, we only consider structure MCMC ( $\mathrm{MC}^{3}$ ) [37] as a comparison. Order MCMC or hybrid DP approaches bound the number of parents and thus often exclude the ground truth graph a priori, especially for scale-free BN structures. Burn-in and thinning for $\mathrm{MC}^{3}$ are chosen to make the wall time comparable with DiBS run on CPUs. In the remainder of the paper, each method uses 30 samples to approximate the posterior over BNs.

Figure 2 summarizes the results for 30 randomly generated BNs with $d=20$ nodes. We find that DiBS+ performs well compared to the other methods, all of which were specifically developed for the marginal inference scenario evaluated here. DiBS+ appears to be preferable to DiBS.
Joint posterior inference When inferring the joint posterior $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$, we can employ a more explicit representation of linear Gaussian BNs, where the conditional distribution parameters are standard Gaussian. Here, DiBS can leverage the Gumbel-softmax estimator in (12) because $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$ is well-defined when $\mathbf{G}$ lies on the interior of $[0,1]^{d \times d}$ (see Appendix E.1). To provide a comparison with DiBS in the absence of an applicable MCMC method, we propose two variants of $\mathrm{MC}^{3}$ as baselines. Metropolis-Hastings $\mathrm{MC}^{3}\left(\mathrm{M}-\mathrm{MC}^{3}\right)$ jointly samples parameters and structures, and Metropolis-within-Gibbs $\mathrm{MC}^{3}\left(\mathrm{G}-\mathrm{MC}^{3}\right)$ alternates in proposing structure and parameters [56]. Moreover, we extend the bootstrap methods by taking the closed-form maximum likelihood estimate [57] as the posterior parameter sample for a given graph inferred using the BGe score ( $\mathrm{BPC}^{*}$ and BGES*), an approach taken in, e.g., causal BN learning [14].
Figure 3 shows the results for $d=20$ nodes, where $\mathbb{E}$-SHD and AUROC are computed by empirically marginalizing out the parameters. We find that DiBS+ is the only considered method that performs well across all of the metrics, often outperforming the baselines by a significant margin. As for marginal posterior inference of linear Gaussian BNs, DiBS+ performs slightly better than DiBS.

### 6.3 Nonlinear Gaussian Bayesian Networks

We also consider joint inference of nonlinear Gaussian BNs where the mean of each local conditional Gaussian is parameterized by a 2-layer dense neural network with five hidden nodes and ReLU activation functions (see Appendix E.1). Since the marginal likelihood does not have a closed form, we are unable to use BPC* and BGES* as a means of comparison. Figure 4 displays the results for $d=20$ variables, where a given BN model has $|\Theta|=2,220$ weights and biases. Analogous to joint inference of linear Gaussian BNs, DiBS and DiBS+ outperform the MCMC baselines across the considered metrics. To the best of our knowledge, this is the first time that such nonlinear Gaussian BN models have been inferred under the Bayesian paradigm, which opens up exciting avenues in the active learning of more sophisticated causal structures.


Figure 4: Joint posterior inference of nonlinear Gaussian BNs with $d=20$ nodes. Here, the mean of the local conditional distribution of each node is parameterized by a 2-layer neural network with five hidden nodes. DiBS and DiBS+ perform favorably across the board, particularly in the graph metrics.

Appendix F complements Sections 6.2 and 6.3 with details on computing time and efficient implementation of DiBS with SVGD, showing that GPU wall times of the above inference experiments lie on the order of seconds or only a few minutes. In addition, Appendix G provides results for $d=50$ variables, where DiBS+ likewise performs favorably when jointly inferring $p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})$. For marginal posterior inference of $p(\mathbf{G} \mid \mathcal{D})$ under the BGe marginal likelihood, DiBS appears to require more Monte Carlo samples to compensate for the high variance of the score function estimator in this high-dimensional setting.

### 6.4 DiBS with SVGD: Additional Analyses and Ablation Studies

Having compared DiBS and its instantiation with SVGD to existing approaches, we finally devote Appendix H to empirically analyzing some properties of the algorithm. One of our key results is that, all other things held equal, substituting our inner product model in (6) with $p_{\alpha}\left(g_{i j}=1 \mid \mathbf{Z}\right)=\sigma_{\alpha}\left(z_{i j}\right)$, where single scalars encode the edge probabilities, results in significantly worse evaluation metrics.

We additionally study the uncertainty quantification in (non)identifiable edge structures and show the effects of reducing the latent dimensionality $k$ or the number of iterations $T$. Our findings suggest that reducing either hyperparameter still allows for competitive posterior approximations and enables trading off posterior inference quality with computational efficiency, e.g., in large-scale applications.

## 7 Application: Inferring Protein Signaling Networks From Cell Data

A widely used benchmark in structure learning is the proteomics data set by Sachs et al. [3]. The data contain $N=7,466$ continuous measurements of $d=11$ proteins involved in human immune system cells as well as an established causal network of their signaling interactions.
We infer both linear and nonlinear Gaussian BNs with Erdős-Rényi graph priors exactly as in Section 6. The AUROC results in Table 1 indicate that the posterior by DiBS under the BGe model provides the most calibrated edge confidence scores. Not penalizing model complexity as much, the marginal BGe posterior of DiBS (DiBS+) averages a high expected number of 39.0 (35.0) edges, compared to 12.7 (14.2) and 12.6 (14.2) for its joint posteriors over linear and nonlinear BNs, respectively. Appendix I provides further details and analyses on this matter. The $\mathbb{E}$-SHD scores show that, among all the methods, DiBS is closest in structure to the consensus network when performing joint inference with nonlinear Gaussian BNs. This further highlights the need for nonlinear conditionals

Table 1: Inference of protein signaling pathways with Gaussian BNs. Metrics are the mean $\pm \mathrm{SD}$ of 30 random restarts.

|  |  | E-SHD | AUROC |
| :--- | :--- | :---: | :---: |
|  | MC $^{3}$ | $34.0 \pm 0.7$ | $0.616 \pm 0.027$ |
|  | BPC | $25.5 \pm 2.3$ | $0.566 \pm 0.020$ |
| $\dagger$ | BGES | $33.7 \pm 1.7$ | $0.641 \pm 0.030$ |
|  | DiBS | $37.4 \pm 0.5$ | $\mathbf{0 . 6 4 7} \pm \mathbf{0 . 0 4 7}$ |
|  | DiBS + | $34.7 \pm 1.5$ | $0.629 \pm 0.045$ |
|  | M-MC $^{3}$ | $37.3 \pm 3.5$ | $0.551 \pm 0.078$ |
| § | G-MC $^{3}$ | $30.5 \pm 3.2$ | $0.527 \pm 0.067$ |
|  | DiBS | $23.4 \pm 0.5$ | $0.598 \pm 0.052$ |
|  | DiBS + | $22.9 \pm 2.7$ | $0.557 \pm 0.052$ |
|  | M-MC $^{3}$ | $25.2 \pm 3.0$ | $0.526 \pm 0.084$ |
| G | G-MC $^{3}$ | $35.1 \pm 3.2$ | $0.540 \pm 0.080$ |
|  | DiBS | $\mathbf{2 2 . 6} \pm \mathbf{0 . 5}$ | $0.577 \pm 0.039$ |
|  | DiBS + | $22.8 \pm 1.9$ | $0.535 \pm 0.041$ |

$\dagger$ Linear Gaussian BN; graph only via BGe marginal lik.
${ }^{\S}$ Linear Gaussian BN; graph and parameters jointly
${ }^{\text {4 }}$ Nonlinear Gaussian BN; graph and parameters jointly and joint inference of the graph and parameters in complex real-world settings.

## 8 Conclusion

We have presented a general, fully differentiable approach to inference of posterior distributions over BNs. Our framework is based on a continuous latent representation of DAGs, whose posterior can be equivalently inferred-without loss of generality and using existing black box inference methods. While we have used SVGD [19] for this purpose, our general approach could also be instantiated with, e.g., gradient-based sampling methods that rely on the score of the target density [58, 59]. This may improve upon the asymptotic runtime of DiBS with SVGD, which scales quadratically in the number of sampled particles. We expect that our end-to-end approach can be extended to handle missing and interventional data as well as to amortized contexts, where rich unstructured data is available.

Broader impact Our work is relevant to any scientific discipline that aims at inferring the (causal) structure of a system or reasoning about the effects of interventions. If algorithms and decisions are grounded in the structural understanding of a data-generating system and take into account the epistemic uncertainty, we expect them to be more robust and have fewer unforeseen side-effects. However, the assumptions allowing for a causal interpretation of DAGs, e.g., the absence of unmeasured confounders, are often untestable and to be taken with care [60], particularly in safety-critical and societally-sensitive applications. Hence, while potential misuse can never be ruled out, our presented method predominantly promises positive societal and scientific impact.

## Acknowledgments and Disclosure of Funding

This project received funding from the Swiss National Science Foundation under NCCR Automation under grant agreement 51NF40 180545, the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program grant agreement no. 815943, and was supported with compute resources by Oracle Cloud Services. This work was also supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B, and by the Machine Learning Cluster of Excellence, EXC number 2064/1 - Project number 390727645. We thank Nicolo Ruggeri and Guillaume Wang for their valuable feedback.

## References

[1] Judea Pearl. Probabilistic reasoning in intelligent systems: networks of plausible inference. 1988.
[2] Dana Pe'er, Aviv Regev, Gal Elidan, and Nir Friedman. Inferring subnetworks from perturbed expression profiles. Bioinformatics, 17:S215-S224, 2001.
[3] Karen Sachs, Omar Perez, Dana Pe'er, Douglas A Lauffenburger, and Garry P Nolan. Causal proteinsignaling networks derived from multiparameter single-cell data. Science, 308(5721):523-529, 2005.
[4] Chikako Van Koten and AR Gray. An application of bayesian network for predicting object-oriented software maintainability. Information and Software Technology, 48(1):59-67, 2006.
[5] Philippe Weber, Gabriela Medina-Oliva, Christophe Simon, and Benoît Iung. Overview on bayesian networks applications for dependability, risk analysis and maintenance areas. Engineering Applications of Artificial Intelligence, 25(4):671-682, 2012.
[6] David Maxwell Chickering. Optimal structure identification with greedy search. J. Mach. Learn. Res., 3: 507-554, March 2003.
[7] Peter Spirtes, Clark N. Glymour, and Richard Scheines. Causation, prediction, and search. Adaptive computation and machine learning. MIT Press, Cambridge, Mass, 2nd ed edition, 2000.
[8] Ioannis Tsamardinos, Laura E Brown, and Constantin F Aliferis. The max-min hill-climbing bayesian network structure learning algorithm. Machine learning, 65(1):31-78, 2006.
[9] Xun Zheng, Bryon Aragam, Pradeep K Ravikumar, and Eric P Xing. DAGs with NO TEARS: Continuous optimization for structure learning. In Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018.
[10] Kevin P Murphy. Active learning of causal bayes net structure. 2001.
[11] Simon Tong and Daphne Koller. Active learning for structure in bayesian networks. In International joint conference on artificial intelligence, volume 17, pages 863-869, 2001.
[12] Hyunghoon Cho, Bonnie Berger, and Jian Peng. Reconstructing causal biological networks through active learning. PloS one, 11(3):e0150611, 2016.
[13] Robert Osazuwa Ness, Karen Sachs, Parag Mallick, and Olga Vitek. A bayesian active learning experimental design for inferring signaling networks. In International Conference on Research in Computational Molecular Biology, pages 134-156. Springer, 2017.
[14] Raj Agrawal, Chandler Squires, Karren Yang, Karthikeyan Shanmugam, and Caroline Uhler. ABCDStrategy: Budgeted experimental design for targeted causal structure discovery. In Proceedings of Machine Learning Research, volume 89, pages 3400-3409. PMLR, 16-18 Apr 2019.
[15] Julius von Kügelgen, Paul K Rubenstein, Bernhard Schölkopf, and Adrian Weller. Optimal experimental design via bayesian optimization: active causal structure learning for gaussian process networks. arXiv preprint arXiv:1910.03962, 2019.
[16] Dan Geiger and David Heckerman. Learning gaussian networks. In Proceedings of the Tenth International Conference on Uncertainty in Artificial Intelligence, UAI'94, page 235-243, San Francisco, CA, USA, 1994.
[17] Dan Geiger and David Heckerman. Parameter priors for directed acyclic graphical models and the characterization of several probability distributions. Ann. Statist., 30(5):1412-1440, October 2002. Publisher: The Institute of Mathematical Statistics.
[18] David Heckerman, Dan Geiger, and David M. Chickering. Learning Bayesian Networks: The Combination of Knowledge and Statistical Data. Machine Learning, 20(3):197-243, September 1995.
[19] Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm. In Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16, page 2378-2386, 2016.
[20] Nir Friedman and Daphne Koller. Being Bayesian About Network Structure. A Bayesian Approach to Structure Discovery in Bayesian Networks. Machine Learning, 50(1):95-125, January 2003.
[21] David Madigan and Adrian E Raftery. Model selection and accounting for model uncertainty in graphical models using occam's window. Journal of the American Statistical Association, 89(428):1535-1546, 1994.
[22] David Madigan, Jonathan Gavrin, and Adrian E Raftery. Eliciting prior information to enhance the predictive performance of bayesian graphical models. Communications in Statistics-Theory and Methods, 24(9):2271-2292, 1995.
[23] R.W. Robinson. Counting labeled acyclic digraphs. New Directions in the Theory of Graphs, 1973.
[24] Yue Yu, Jie Chen, Tian Gao, and Mo Yu. DAG-GNN: DAG structure learning with graph neural networks. In Proceedings of the 36th International Conference on Machine Learning, volume 97 of Proceedings of Machine Learning Research, pages 7154-7163. PMLR, 09-15 Jun 2019.
[25] Sébastien Lachapelle, Philippe Brouillard, Tristan Deleu, and Simon Lacoste-Julien. Gradient-based neural dag learning. In International Conference on Learning Representations, 2020.
[26] Ignavier Ng, Shengyu Zhu, Zhitang Chen, and Zhuangyan Fang. A graph autoencoder approach to causal structure learning. In NeurIPS 2019 Workshop "Do the right thing": machine learning and causal inference for improved decision making, 2019.
[27] Xun Zheng, Chen Dan, Bryon Aragam, Pradeep Ravikumar, and Eric Xing. Learning sparse nonparametric dags. In International Conference on Artificial Intelligence and Statistics, pages 3414-3425. PMLR, 2020.
[28] Roxana Pamfil, Nisara Sriwattanaworachai, Shaan Desai, Philip Pilgerstorfer, Konstantinos Georgatzis, Paul Beaumont, and Bryon Aragam. Dynotears: Structure learning from time-series data. In Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics, volume 108 of Proceedings of Machine Learning Research, pages 1595-1605. PMLR, 2020.
[29] Mengyue Yang, Furui Liu, Zhitang Chen, Xinwei Shen, Jianye Hao, and Jun Wang. Causalvae: Disentangled representation learning via neural structural causal models. arXiv preprint arXiv: 2004.08697, 2020.
[30] Muhan Zhang, Shali Jiang, Zhicheng Cui, Roman Garnett, and Yixin Chen. D-vae: A variational autoencoder for directed acyclic graphs. In Advances in Neural Information Processing Systems, volume 32, 2019.
[31] Nan Rosemary Ke, Olexa Bilaniuk, Anirudh Goyal, Stefan Bauer, Hugo Larochelle, Chris Pal, and Yoshua Bengio. Learning neural causal models from unknown interventions. arXiv preprint arXiv: 1910.01075, 2019.
[32] Ignavier Ng, Zhuangyan Fang, Shengyu Zhu, Zhitang Chen, and Jun Wang. Masked gradient-based causal structure learning. arXiv preprint arXiv:1910.08527, 2019.
[33] Philippe Brouillard, Sébastien Lachapelle, Alexandre Lacoste, Simon Lacoste-Julien, and Alexandre Drouin. Differentiable causal discovery from interventional data. arXiv preprint arXiv:2007.01754, 2020.
[34] Dennis Wei, Tian Gao, and Yue Yu. Dags with no fears: A closer look at continuous optimization for learning bayesian networks. Advances in Neural Information Processing Systems, 2020.
[35] Ignavier Ng, Sébastien Lachapelle, Nan Rosemary Ke, and Simon Lacoste-Julien. On the convergence of continuous constrained optimization for structure learning. arXiv preprint arXiv:2011.11150, 2020.
[36] David Madigan, Jeremy York, and Denis Allard. Bayesian graphical models for discrete data. International Statistical Review/Revue Internationale de Statistique, pages 215-232, 1995.
[37] Paolo Giudici and Robert Castelo. Improving markov chain monte carlo model search for data mining. Machine learning, 50(1-2):127-158, 2003.
[38] Daniel Eaton and Kevin Murphy. Bayesian structure learning using dynamic programming and MCMC. In Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI), 2007.
[39] Marco Grzegorczyk and Dirk Husmeier. Improving the structure MCMC sampler for Bayesian networks by introducing a new edge reversal move. Machine Learning, 71(2):265, April 2008.
[40] Jack Kuipers and Giusi Moffa. Partition mcmc for inference on acyclic digraphs. Journal of the American Statistical Association, 112(517):282-299, 2017.
[41] Byron Ellis and Wing Hung Wong. Learning causal bayesian network structures from experimental data. Journal of the American Statistical Association, 103(482):778-789, 2008.
[42] N. Friedman, M. Goldszmidt, and A. J. Wyner. Data analysis with Bayesian networks: A bootstrap approach. In Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, 1999.
[43] Mikko Koivisto and Kismat Sood. Exact bayesian structure discovery in bayesian networks. Journal of Machine Learning Research, 5(May):549-573, 2004.
[44] Mikko Koivisto. Advances in exact bayesian structure discovery in bayesian networks. In Proceedings of the Twenty-Second Conference on Uncertainty in Artificial Intelligence, page 241-248, 2006.
[45] Marina Meila and Tommi S Jaakkola. Tractable bayesian learning of tree belief networks. Statistics and Computing, (16):77-92, 2006.
[46] Denver Dash and Gregory F Cooper. Model averaging for prediction with discrete bayesian networks. Journal of Machine Learning Research, 5(Sep):1177-1203, 2004.
[47] Thomas N Kipf and Max Welling. Variational graph auto-encoders. arXiv preprint arXiv:1611.07308, 2016.
[48] Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In Proceedings International Conference on Learning Representations, 2017.
[49] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparametrization with gumbel-softmax. In Proceedings International Conference on Learning Representations 2017, April 2017.
[50] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine learning, 8(3-4):229-256, 1992.
[51] Shakir Mohamed, Mihaela Rosca, Michael Figurnov, and Andriy Mnih. Monte carlo gradient estimation in machine learning. Journal of Machine Learning Research, 21(132):1-62, 2020.
[52] Ignavier Ng, AmirEmad Ghassami, and Kun Zhang. On the role of sparsity and dag constraints for learning linear dags. Advances in Neural Information Processing Systems, 2020.
[53] P. Erdős and A Rényi. On random graphs. Publicationes Mathematicae, 6:290-297, 1959.
[54] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. Science, 286(5439): 509-512, 1999.
[55] Steen A Andersson, David Madigan, Michael D Perlman, et al. A characterization of markov equivalence classes for acyclic digraphs. Annals of statistics, 25(2):505-541, 1997.
[56] W Keith Hastings. Monte carlo sampling methods using markov chains and their applications. 1970.
[57] Alain Hauser and Peter Bühlmann. Jointly interventional and observational data: estimation of interventional markov equivalence classes of directed acyclic graphs. Journal of the Royal Statistical Society: Series B: Statistical Methodology, pages 291-318, 2015.
[58] Radford M Neal et al. Mcmc using hamiltonian dynamics. Handbook of Markov chain Monte Carlo, 2 (11):2, 2011.
[59] Max Welling and Yee Whye Teh. Bayesian learning via stochastic gradient langevin dynamics. In Proceedings of the 28th International Conference on International Conference on Machine Learning, ICML'11, page 681-688, 2011.
[60] A. Philip Dawid. Beware of the DAG! In Proceedings of Workshop on Causality: Objectives and Assessment at NIPS 2008, pages 59-86, 2010.
[61] Jonas Peters and Peter Bühlmann. Identifiability of gaussian structural equation models with equal error variances. Biometrika, 101(1):219-228, 2014.
[62] Jack Kuipers, Giusi Moffa, and David Heckerman. Addendum on the scoring of Gaussian directed acyclic graphical models. Ann. Statist., 42(4):1689-1691, August 2014. Publisher: The Institute of Mathematical Statistics.
[63] Diviyan Kalainathan, Olivier Goudet, Isabelle Guyon, David Lopez-Paz, and Michèle Sebag. Structural agnostic modeling: Adversarial learning of causal graphs. arXiv preprint arXiv:1803.04929, 2018.
[64] Chris Sherlock, Paul Fearnhead, Gareth O Roberts, et al. The random walk metropolis: linking theory and practice through a case study. Statistical Science, 25(2):172-190, 2010.
[65] Dorit Dor and Michael Tarsi. A simple algorithm to construct a consistent extension of a partially oriented graph. 1992.
[66] Yangbo He, Jinzhu Jia, and Bin Yu. Counting and exploring sizes of markov equivalence classes of directed acyclic graphs. The Journal of Machine Learning Research, 16(1):2589-2609, 2015.
[67] Diviyan Kalainathan and Olivier Goudet. Causal discovery toolbox: Uncover causal relationships in python. arXiv preprint arXiv:1903.02278, 2019. URL
https://github.com/FenTechSolutions/CausalDiscoveryToolbox.
[68] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL http://github. com/google/jax.
[69] Yoshua Bengio, Tristan Deleu, Nasim Rahaman, Nan Rosemary Ke, Sebastien Lachapelle, Olexa Bilaniuk, Anirudh Goyal, and Christopher Pal. A meta-transfer objective for learning to disentangle causal mechanisms. In International Conference on Learning Representations, 2020.

## A Proofs of the Main Results

## A. 1 Proposition 1

Proof For ease of understanding, we recall that the generative model in (5) factorizes the joint distribution as $p(\mathbf{Z}, \mathbf{G}, \boldsymbol{\Theta}, \mathcal{D})=p(\mathbf{Z}) p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})$. First, let us consider case (a), the setting where the marginal likelihood $p(\mathcal{D} \mid \mathbf{G})$ can be computed in closed form. We get

$$
\begin{align*}
\mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})} f(\mathbf{G}) & =\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathcal{D}) f(\mathbf{G})  \tag{A.1}\\
& =\sum_{\mathbf{G}} \frac{p(\mathbf{G}, \mathcal{D}) f(\mathbf{G})}{p(\mathcal{D})}  \tag{A.2}\\
& =\sum_{\mathbf{G}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z}, \mathbf{G}, \mathcal{D}) f(\mathbf{G})}{p(\mathcal{D})} d \mathbf{Z}  \tag{A.3}\\
& =\sum_{\mathbf{G}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z}) p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G}) f(\mathbf{G})}{p(\mathcal{D})} d \mathbf{Z} \tag{A.4}
\end{align*}
$$

by the generative model in (5)

$$
\begin{align*}
& =\sum_{\mathbf{G}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z} \mid \mathcal{D}) p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G}) f(\mathbf{G})}{p(\mathcal{D} \mid \mathbf{Z})} d \mathbf{Z}  \tag{A.5}\\
& \quad \text { since } p(\mathbf{Z} \mid \mathcal{D})=\frac{p(\mathbf{Z}) p(\mathcal{D} \mid \mathbf{Z})}{p(\mathcal{D})} \\
& =\int_{\mathbf{Z}} p(\mathbf{Z} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G}) f(\mathbf{G})}{p(\mathcal{D} \mid \mathbf{Z})} d \mathbf{Z} \quad \text { rearranging }  \tag{A.6}\\
& =\int_{\mathbf{Z}} p(\mathbf{Z} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G}) f(\mathbf{G})}{\sum_{\mathbf{G}} p(\mathbf{G}, \mathcal{D} \mid \mathbf{Z})} d \mathbf{Z} \tag{A.7}
\end{align*}
$$

by the law of total probability
$=\int_{\mathbf{Z}} p(\mathbf{Z} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G}) f(\mathbf{G})}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G})} d \mathbf{Z}$
expanding $p(\mathbf{G}, \mathcal{D} \mid \mathbf{Z})$ by the generative model in (5)
$=\mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}) p(\mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]}\right]$
as desired for (a).
Finally, let us consider (b), the general case. The derivation essentially follows the same ideas as for (a) but does not marginalize out $\Theta$.

$$
\begin{align*}
& \mathbb{E}_{p(\mathbf{G}, \boldsymbol{\Theta} \mid D)} f(\mathbf{G}, \boldsymbol{\Theta})  \tag{A.10}\\
& =\sum_{\mathbf{G}} \int_{\boldsymbol{\Theta}} p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D}) f(\mathbf{G}, \boldsymbol{\Theta}) d \boldsymbol{\Theta}  \tag{A.11}\\
& =\sum_{\mathbf{G}} \int_{\boldsymbol{\Theta}} \frac{p(\mathbf{G}, \boldsymbol{\Theta}, \mathcal{D}) f(\mathbf{G}, \boldsymbol{\Theta})}{p(\mathcal{D})} d \boldsymbol{\Theta}  \tag{A.12}\\
& =\sum_{\mathbf{G}} \int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z}, \mathbf{G}, \boldsymbol{\Theta}, \mathcal{D}) f(\mathbf{G}, \boldsymbol{\Theta})}{p(\mathcal{D})} d \mathbf{Z} d \boldsymbol{\Theta} \tag{A.13}
\end{align*}
$$

$$
=\sum_{\mathbf{G}} \int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z}) p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) f(\mathbf{G}, \boldsymbol{\Theta})}{p(\mathcal{D})} d \mathbf{Z} d \boldsymbol{\Theta}
$$

by the generative model in (5)

$$
\begin{align*}
& =\sum_{\mathbf{G}} \int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} \frac{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D}) p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) f(\mathbf{G}, \boldsymbol{\Theta})}{p(\mathcal{D}, \boldsymbol{\Theta} \mid \mathbf{Z})} d \mathbf{Z} d \mathbf{\Theta}  \tag{A.15}\\
& \text { since } p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})=\frac{p(\mathbf{Z}) p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})}{p(\mathcal{D})} \\
& =\int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) f(\mathbf{G}, \boldsymbol{\Theta})}{p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})} d \mathbf{Z} d \mathbf{\Theta} \tag{A.16}
\end{align*}
$$

rearranging

$$
\begin{equation*}
=\int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) f(\mathbf{G}, \boldsymbol{\Theta})}{\sum_{\mathbf{G}} p(\mathbf{G}, \boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})} d \mathbf{Z} d \mathbf{\Theta} \tag{A.17}
\end{equation*}
$$

by the law of total probability
$=\int_{\boldsymbol{\Theta}} \int_{\mathbf{Z}} p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D}) \frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta}) f(\mathbf{G}, \boldsymbol{\Theta})}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})} d \mathbf{Z} d \boldsymbol{\Theta}$
expanding $p(\mathbf{G}, \boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})$ by the generative model in (5)
$=\mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}, \boldsymbol{\Theta}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}\right]$
which is the statement in (b).

## A. 2 Proposition 2

Proof We will derive the gradients of the unnormalized posterior since

$$
\begin{equation*}
\nabla_{\mathbf{Z}} \log p(\mathbf{Z} \mid \mathcal{D})=\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \mathcal{D})-\nabla_{\mathbf{Z}} \log p(\mathcal{D})=\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \mathcal{D}) \tag{A.20}
\end{equation*}
$$

and analogously for the other two expressions. Through straightforward manipulation and using the identity $\nabla_{\mathbf{x}} \log f(\mathbf{x})=\nabla_{\mathbf{x}} f(\mathbf{x}) / f(\mathbf{x})$, we obtain

$$
\begin{align*}
\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \mathcal{D}) & =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\nabla_{\mathbf{Z}} \log p(\mathcal{D} \mid \mathbf{Z})  \tag{A.21}\\
& =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} p(\mathcal{D} \mid \mathbf{Z})}{p(\mathcal{D} \mid \mathbf{Z})}  \tag{A.22}\\
& =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}}\left[\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G})\right]}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathcal{D} \mid \mathbf{G})}  \tag{A.23}\\
& =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]} \tag{A.24}
\end{align*}
$$

Analogously, we get

$$
\begin{align*}
\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \boldsymbol{\Theta}, \mathcal{D}) & =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\nabla_{\mathbf{Z}} \log p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})  \tag{A.25}\\
& =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})}{p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})}  \tag{A.26}\\
& =\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}}\left[\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})\right]}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})} \tag{A.27}
\end{align*}
$$

$$
\begin{equation*}
=\nabla_{\mathbf{Z}} \log p(\mathbf{Z})+\frac{\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})]} \tag{A.28}
\end{equation*}
$$

Lastly, using the same ideas as above, we arrive at

$$
\begin{align*}
\nabla_{\boldsymbol{\Theta}} \log p(\mathbf{Z}, \boldsymbol{\Theta}, \mathcal{D}) & =\nabla_{\boldsymbol{\Theta}} \log p(\mathbf{Z})+\nabla_{\boldsymbol{\Theta}} \log p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})  \tag{A.29}\\
& =\frac{\nabla_{\boldsymbol{\Theta}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})}{p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{Z})}  \tag{A.30}\\
& =\frac{\nabla_{\boldsymbol{\Theta}}\left[\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})\right]}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})}  \tag{A.31}\\
& =\frac{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) \nabla_{\boldsymbol{\Theta}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})}{\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathbf{Z}) p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})}  \tag{A.32}\\
& =\frac{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}\left[\nabla_{\boldsymbol{\Theta}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})\right]}{\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[p(\mathbf{\Theta}, \mathcal{D} \mid \mathbf{G})]} \tag{A.33}
\end{align*}
$$

In the above, without any additional factor modeling a prior belief over graphs, the score of the latent prior $p_{\beta}(\mathbf{Z})$ as defined in (8) is given by

$$
\begin{equation*}
\nabla_{\mathbf{Z}} \log p_{\beta}(\mathbf{Z})=-\beta \nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[h(\mathbf{G})]-\frac{1}{\sigma_{z}^{2}} \mathbf{Z} \tag{A.34}
\end{equation*}
$$

Practical considerations Estimating expectations of the form $\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})]$ with Monte Carlo sampling can be numerically challenging when $f$ are probability densities and thus often close to zero. In practice, we recommend the log-sum-exp trick for applying Proposition 2. Let us define

$$
\begin{equation*}
\underset{m=1}{M} \sum_{m} \mathrm{E}\left\{x^{(m)}\right\}:=\log \left(\sum_{m=1}^{M} \exp \left(x^{(m)}\right)\right) \tag{A.35}
\end{equation*}
$$

For $M$ Monte Carlo samples $\mathbf{G}^{(m)} \sim p(\mathbf{G} \mid \mathbf{Z})$, we can rewrite the estimator for the expectation as

$$
\begin{align*}
\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})] & \approx \frac{1}{M} \sum_{m=1}^{M} f\left(\mathbf{G}^{(m)}\right)=\exp \left(\log \left(\sum_{m=1}^{M} f\left(\mathbf{G}^{(m)}\right)\right)-\log M\right)  \tag{A.36}\\
& =\exp \left(\underset{m=1}{\operatorname{L} E}\left\{\log f\left(\mathbf{G}^{(m)}\right)\right\}-\log M\right) \tag{A.37}
\end{align*}
$$

Computing LEE can be made numerically stable by subtracting and adding $\max _{m}\left\{x^{(m)}\right\}$ before and after applying $L \Sigma E$ to $\left\{x^{(m)}\right\}$, respectively. Stable $L \Sigma E$ can be extended to handle negative-valued $f$ inside the expectation, e.g., for the gradient of $f$, and to the ratios of expectations in Proposition 2.

## A. 3 Equations (16) and (17)

Proof The sigmoid function converges to the unit step function, i.e. $\sigma_{\alpha}(x) \rightarrow \mathbb{1}[x>0]$ as $\alpha \rightarrow \infty$. Hence, the edge probabilities $\mathbf{G}_{\alpha}(\mathbf{Z})$ defined in (7) converge to a (binary) matrix $\mathbf{G}$ as $\alpha \rightarrow \infty$. Extending the notation of (7), we will denote this single limiting graph implied by $\mathbf{Z}$ as $\mathbf{G}_{\infty}(\mathbf{Z})$ where

$$
\mathbf{G}_{\infty}(\mathbf{Z})_{i j}:= \begin{cases}1 & \text { if } \mathbf{u}_{i}^{\top} \mathbf{v}_{j}>0 \text { and } i \neq j \\ 0 & \text { otherwise }\end{cases}
$$

The above implies that when the temperature parameter $\alpha$ is annealed to $\infty$, the probability mass function and correspondingly the expectation simplify:

$$
\begin{align*}
\text { As } \alpha \rightarrow \infty: & p_{\alpha}(\mathbf{G} \mid \mathbf{Z}) \rightarrow \mathbb{1}\left[\mathbf{G}=\mathbf{G}_{\infty}(\mathbf{Z})\right] \\
& \mathbb{E}_{p_{\alpha}(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})] \rightarrow f\left(\mathbf{G}_{\infty}(\mathbf{Z})\right) \tag{A.38}
\end{align*}
$$

Again, let us first consider case (a). Starting with Proposition 1(a) in the first step, we can use the above insight to simplify the inner expectations:

$$
\begin{equation*}
\mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})}[f(\mathbf{G})]=\mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p_{\alpha}(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}) p(\mathcal{D} \mid \mathbf{G})]}{\mathbb{E}_{p_{\alpha}(\mathbf{G} \mid \mathbf{Z})}[p(\mathcal{D} \mid \mathbf{G})]}\right] \tag{A.39}
\end{equation*}
$$

$$
\begin{align*}
& \xrightarrow{\alpha \rightarrow \infty}  \tag{A.40}\\
& \quad \mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[\frac{f\left(\mathbf{G}_{\infty}(\mathbf{Z})\right) p\left(\mathcal{D} \mid \mathbf{G}_{\infty}(\mathbf{Z})\right)}{p\left(\mathcal{D} \mid \mathbf{G}_{\infty}(\mathbf{Z})\right)}\right]  \tag{A.41}\\
& \quad=\mathbb{E}_{p(\mathbf{Z} \mid \mathcal{D})}\left[f\left(\mathbf{G}_{\infty}(\mathbf{Z})\right)\right]
\end{align*}
$$

Analogously, we get for the general case (b):

$$
\begin{align*}
\mathbb{E}_{p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})}[f(\mathbf{G}, \boldsymbol{\Theta})] & =\mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[\frac{\mathbb{E}_{p_{\alpha}(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G}, \boldsymbol{\Theta}) p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}{\mathbb{E}_{p_{\alpha}(\mathbf{G} \mid \mathbf{Z})}[p(\boldsymbol{\Theta} \mid \mathbf{G}) p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})]}\right]  \tag{A.42}\\
\xrightarrow{\alpha \rightarrow \infty} & \mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[\frac{f\left(\mathbf{G}_{\infty}(\mathbf{Z}), \boldsymbol{\Theta}\right) p\left(\boldsymbol{\Theta} \mid \mathbf{G}_{\infty}(\mathbf{Z})\right) p\left(\mathcal{D} \mid \mathbf{G}_{\infty}(\mathbf{Z}), \boldsymbol{\Theta}\right)}{p\left(\boldsymbol{\Theta} \mid \mathbf{G}_{\infty}(\mathbf{Z})\right) p\left(\mathcal{D} \mid \mathbf{G}_{\infty}(\mathbf{Z}), \boldsymbol{\Theta}\right)}\right]  \tag{A.43}\\
& =\mathbb{E}_{p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})}\left[f\left(\mathbf{G}_{\infty}(\mathbf{Z}), \boldsymbol{\Theta}\right)\right] \tag{A.44}
\end{align*}
$$

## B Gradient Estimation for Bayesian Inference

To derive the expressions for the gradient estimators in (12) and (14) for both the marginal likelihood and the likelihood cases, we will use a generic function $f(\mathbf{G})$ as a placeholder for either $p(\mathcal{D} \mid \mathbf{G})$ or $p(\mathcal{D} \mid \mathbf{G}, \boldsymbol{\Theta})$, since the results hold for general densities $f(\mathbf{G})$.

## B. 1 Gumbel-Softmax Estimator for the Likelihood Gradient

In general, for a Bernoulli random variable $X$ with $p(X=1)=q$, it holds that

$$
\begin{equation*}
X \stackrel{d}{=} \mathbb{1}\left[G_{1}+\log q>G_{0}+\log (1-q)\right] \tag{B.1}
\end{equation*}
$$

when $G_{0}, G_{1} \sim \operatorname{Gumbel}(0,1)$. This is the Gumbel-max trick. Since the unit step function $\mathbb{1}[\cdot]$ does not have an informative gradient, Maddison et al. [48] and Jang et al. [49] have proposed to use the sigmoid function $\sigma_{\tau}(\cdot)$ with parameter $\tau$ as a soft relaxation of $\mathbb{1}[\cdot]$.

Using this so-called Gumbel-softmax trick, we can reparameterize the entries of $\mathbf{G}$ under the graph model in (6). Starting from the Gumbel-max equality in (B.1), we rearrange the inequality inside the indicator into the form " $>0$ " and apply the sigmoid relaxation with parameter $\tau$. We obtain the following soft relaxation for each entry of $\mathbf{G}$ :

$$
\begin{align*}
g_{i j} & \approx \sigma_{\tau}\left(G_{1}-G_{0}+\log \sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)-\log \left(1-\sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)\right)\right)  \tag{B.2}\\
& =\sigma_{\tau}\left(L+\log \left(\frac{\sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)}{\sigma_{\alpha}\left(-\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)}\right)\right)  \tag{B.3}\\
& =\sigma_{\tau}\left(L+\log \left(\frac{\exp \left(\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)}{\exp \left(\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)+1} \frac{\exp \left(\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)+1}{1}\right)\right)  \tag{B.4}\\
& =\sigma_{\tau}\left(L+\log \left(\exp \left(\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)\right)\right)  \tag{B.5}\\
& =\sigma_{\tau}\left(L+\alpha \mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right) \tag{B.6}
\end{align*}
$$

where $L \sim \operatorname{Logistic}(0,1)$ since $L \stackrel{d}{=} G_{1}-G_{0}$ when $G_{0}, G_{1} \sim \operatorname{Gumbel}(0,1)$. For $i=j$, we set $g_{i j}:=0$ by default in accordance with the graph model in (6). Since this allows us to separate the randomness in sampling from the distribution $p(\mathbf{G} \mid \mathbf{Z})$ from the values of $\mathbf{Z}$, we can move the gradient operator inside the expectation and obtain the estimator given in (12):

$$
\begin{align*}
\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})] & \approx \mathbb{E}_{p(\mathbf{L})}\left[\nabla_{\mathbf{Z}} f\left(\mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})\right)\right] \\
& =\mathbb{E}_{p(\mathbf{L})}\left[\left.\nabla_{\mathbf{G}} f(\mathbf{G})\right|_{\mathbf{G}=\mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})} \cdot \nabla_{\mathbf{Z}} \mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})\right] \tag{B.7}
\end{align*}
$$

While the reparameterization trick generally provides a lower variance estimate of the gradient, the form in (12) is biased when $\tau<\infty$ because we use a soft relaxation of the true distribution. In addition, the estimator in (12) requires that $\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})$ or $\nabla_{\mathbf{G}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ is available, depending
on the inference task. In case $p(\mathcal{D} \mid \mathbf{G})$ or $p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ is only defined for discrete $\mathbf{G}$, it is possible to evaluate $\nabla_{\mathbf{G}} p(\mathcal{D} \mid \mathbf{G})$ or $\nabla_{\mathbf{G}} p(\boldsymbol{\Theta}, \mathcal{D} \mid \mathbf{G})$ using hard Gumbel-max samples of $\mathbf{G}$ (i.e., with $\tau=\infty$ ). As before, however, one would use soft Gumbel-softmax samples in $\nabla_{\mathbf{Z}} \mathbf{G}_{\tau}(\mathbf{L}, \mathbf{Z})$ to obtain an informative gradient. Lastly, we can use this estimator to approximate the score of the latent prior $\nabla_{\mathbf{Z}} \log p(\mathbf{Z})$ given in (8) because the acyclicity constraint $h(\mathbf{G})$ is differentiable with respect to $\mathbf{G}$. In practice, the log-sum-exp trick described in Section A. 2 as well as the score function identity $\nabla_{\mathbf{x}} f(\mathbf{x})=f(\mathbf{x}) \nabla_{\mathbf{x}} \log f(\mathbf{x})$ should be used for numerically stable computation of the estimator.

## B. 2 Score Function Estimator for the Likelihood Gradient

To arrive at the estimator in (14), we expand the gradient expression as

$$
\begin{align*}
\nabla_{\mathbf{Z}} \mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}[f(\mathbf{G})] & =\sum_{\mathbf{G}} f(\mathbf{G}) \nabla_{\mathbf{Z}} p(\mathbf{G} \mid \mathbf{Z})=\sum_{\mathbf{G}} f(\mathbf{G}) p(\mathbf{G} \mid \mathbf{Z}) \nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})  \tag{B.8}\\
& =\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}\left[f(\mathbf{G}) \nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})\right] \tag{B.9}
\end{align*}
$$

where (B.8) uses the fact that $\nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})=\nabla_{\mathbf{Z}} p(\mathbf{G} \mid \mathbf{Z}) / p(\mathbf{G} \mid \mathbf{Z})$. Finally, we recall the well-known property of the score function that $\mathbb{E}_{p(\mathbf{G} \mid \mathbf{Z})}\left[\nabla_{\mathbf{Z}} \log p(\mathbf{G} \mid \mathbf{Z})\right]=\mathbf{0}$. Due to this, for any constant $b$ as written in (14), the estimator is unbiased because the additional term involving $b$ has zero expectation. The constant can be used to reduce the variance of the Monte Carlo estimator [51]. In our experiments, we always use $b=0$.

## C Background: Stein Variational Gradient Descent

This section describes Stein variational gradient descent (SVGD) by Liu and Wang [19]. The overview is meant as supplementary material for Section 5, where we propose to use SVGD for inferring the DiBS posteriors $p(\mathbf{Z} \mid \mathcal{D})$ and $p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})$. In contrast to sampling-based MCMC or optimizationbased variational inference methods, SVGD iteratively transports a fixed set of particles to closely match a target distribution, akin to the gradient descent algorithm in optimization. We refer the reader to Liu and Wang [19] for additional details.
Let $p(\mathbf{x})$ with $\mathbf{x} \in \mathcal{X}$ be a differentiable density that we want to sample from, e.g., to estimate an expectation. Starting from a smooth reference density $q(\mathbf{x})$, SVGD aims to find a one-to-one transform $\mathbf{t}: \mathcal{X} \mapsto \mathcal{X}$ such that the transformed density $q_{[\mathbf{t}]}(\widetilde{\mathbf{x}})$ with $\widetilde{\mathbf{x}}:=\mathbf{t}(\mathbf{x})$ minimizes the KL-divergence to $p$. In particular, Liu and Wang [19] propose to use the incremental transform

$$
\begin{equation*}
\mathbf{t}(\mathbf{x})=\mathbf{x}+\eta \boldsymbol{\phi}(\mathbf{x}) \tag{C.1}
\end{equation*}
$$

When $|\eta|$ is sufficiently small, the Jacobian of $\mathbf{t}$ has full rank and $\mathbf{t}$ is one-to-one. The key result by Liu and Wang [19] links the incremental transform $t$ in (C.1) to prior work on reproducing kernel Hilbert spaces (RKHSs). The authors show that if $\phi$ lies in the unit ball of the RKHS induced by a kernel $k$, then the transform $\mathbf{t}$ maximizing the descent on the KL divergence from $q_{[\mathbf{t}]}$ to $p$ uses an incremental update $\phi$ proportional to

$$
\begin{equation*}
\boldsymbol{\phi}_{q, p}^{*}(\cdot)=\mathbb{E}_{q(\mathbf{x})}\left[k(\mathbf{x}, \cdot) \nabla_{\mathbf{x}} \log p(\mathbf{x})^{\top}+\nabla_{\mathbf{x}} k(\mathbf{x}, \cdot)\right] \tag{C.2}
\end{equation*}
$$

This suggests an iterative procedure of repeatedly applying the update of (C.1) with $\phi=\phi_{q, p}^{*}(\cdot)$ from (C.2) to a finite set of randomly initialized particles $\left\{\mathbf{x}^{(m)}\right\}_{m=1}^{M}$. At each iteration $t$, the $m$-th particle $\mathbf{x}^{(m)}$ is then deterministically updated according to:

$$
\begin{align*}
& \mathbf{x}_{t+1}^{(m)} \leftarrow \mathbf{x}_{t}^{(m)}+\eta_{t} \phi\left(\mathbf{x}_{t}^{(m)}\right) \\
& \text { where } \phi(\mathbf{x})=\frac{1}{M} \sum_{k=1}^{M}\left[k\left(\mathbf{x}_{t}^{(k)}, \mathbf{x}\right) \nabla_{\mathbf{x}_{t}^{(k)}} \log p\left(\mathbf{x}_{t}^{(k)}\right)+\nabla_{\mathbf{x}_{t}^{(k)}} k\left(\mathbf{x}_{t}^{(k)}, \mathbf{x}\right)\right] \tag{C.3}
\end{align*}
$$

For sufficiently small step sizes $\eta_{t}$, the sequence of particles eventually converges, in which case the transform $t$ reduces to the identity mapping. The particle update in (C.3) consists of a gradient ascent term driving the particles to high-density regions, and a term involving $\nabla_{\mathbf{x}} k(\mathbf{x}, \cdot)$ that acts as a repulsive force between particles, preventing them from collapsing into the modes of $p(\mathbf{x})$.

## D General Algorithm

```
Algorithm 2 DiBS with SVGD [19] for inference of \(p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})\)
    Input: Initial latent and parameter particles \(\left\{\left(\mathbf{Z}_{0}^{(m)}, \mathbf{\Theta}_{0}^{(m)}\right)\right\}_{m=1}^{M}\), kernel \(k\), schedules for \(\eta_{t}, \alpha_{t}, \beta_{t}\)
    Output: Set of graph and parameter particles \(\left\{\left(\mathbf{G}^{(m)}, \mathbf{\Theta}^{(m)}\right)\right\}_{m=1}^{M}\) approximating \(p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D})\)
    Incorporate prior belief of \(p(\mathbf{G})\) into \(p(\mathbf{Z}) \quad \triangleright\) See Section 4.2
    for iteration \(t=0\) to \(T-1\) do
        Estimate score \(\nabla_{\mathbf{Z}} \log p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})\) given in (10) for each \(\mathbf{Z}_{t}^{(m)} \quad \triangleright\) See (12) and (14)
        Estimate score \(\nabla_{\boldsymbol{\Theta}} \log p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})\) given in (11) for each \(\boldsymbol{\Theta}_{t}^{(m)}\)
        for particle \(m=1\) to \(M\) do
            \(\mathbf{Z}_{t+1}^{(m)} \leftarrow \mathbf{Z}_{t}^{(m)}+\eta_{t} \boldsymbol{\phi}_{t}^{\mathbf{Z}}\left(\mathbf{Z}_{t}^{(m)}, \mathbf{\Theta}_{t}^{(m)}\right) \quad \triangleright\) SVGD step
            where \(\boldsymbol{\phi}_{t}^{\mathbf{Z}}(\cdot, \cdot):=\frac{1}{M} \sum_{k=1}^{M}\left[k\left(\left(\mathbf{Z}_{t}^{(k)}, \mathbf{\Theta}_{t}^{(k)}\right),(\cdot, \cdot)\right) \nabla_{\mathbf{Z}_{t}^{(k)}} \log p\left(\mathbf{Z}_{t}^{(k)}, \mathbf{\Theta}_{t}^{(k)} \mid \mathcal{D}\right)\right.\)
                        \(\left.+\nabla_{\mathbf{z}_{t}^{(k)}} k\left(\left(\mathbf{Z}_{t}^{(k)}, \boldsymbol{\Theta}_{t}^{(k)}\right),(\cdot, \cdot)\right)\right]\)
            \(\boldsymbol{\Theta}_{t+1}^{(m)} \leftarrow \boldsymbol{\Theta}_{t}^{(m)}+\eta_{t} \boldsymbol{\phi}_{t}^{\boldsymbol{\Theta}}\left(\mathbf{Z}_{t}^{(m)}, \boldsymbol{\Theta}_{t}^{(m)}\right)\)
                where \(\boldsymbol{\phi}_{t}^{\boldsymbol{\Theta}}(\cdot, \cdot)\) is analogous to \(\boldsymbol{\phi}_{t}^{\mathbf{Z}}(\cdot, \cdot)\) but using gradients \(\nabla_{\boldsymbol{\Theta}_{t}^{(k)}}\) instead of \(\nabla_{\mathbf{z}_{t}^{(k)}}\)
    return \(\left\{\left(\mathbf{G}_{\infty}\left(\mathbf{Z}_{T}^{(m)}\right), \mathbf{\Theta}_{T}^{(m)}\right)\right\}_{m=1}^{M} \quad \triangleright \operatorname{See}(16)\) and (17)
```


## E Experimental Details

## E. 1 Gaussian Bayesian Networks

In our experiments, we consider Bayesian networks with Gaussian local conditional distributions of each variable given its parents. For both linear or nonlinear Gaussian BNs, which will be defined presently, the generative model for synthetic data simulation as well as the parameter prior used for joint inference are set to standard Gaussian distributions. We fix the observation noise to $\sigma^{2}=0.1$ for all nodes both during synthetic data generation and joint posterior inference, rendering the causal structure fully identifiable [61].

Linear Analogous to linear regression, linear Gaussian BNs model the mean of a given variable as a linear function of its parents:

$$
\begin{align*}
p(\mathbf{x} \mid \mathbf{G}, \boldsymbol{\Theta}) & =\prod_{i=1}^{d} \mathcal{N}\left(x_{i} ; \boldsymbol{\theta}_{i}^{\top} \mathbf{x}_{\mathrm{pa}(i)}, \sigma^{2}\right)  \tag{E.1}\\
\text { or } \quad p(\mathbf{x} \mid \mathbf{G}, \boldsymbol{\Theta}) & =\mathcal{N}\left(\mathbf{x} ;(\mathbf{G} \circ \boldsymbol{\Theta})^{\top} \mathbf{x}, \sigma^{2} \mathbf{I}\right)
\end{align*}
$$

where " $\circ$ " denotes elementwise multiplication. In our experiments, DiBS uses the second parameterization in (E.1) to allow for a constant dimensionality of the conditional distribution parameters $\Theta$ and make the likelihood well-defined for the Gumbel-softmax estimator in (12).

When inferring the marginal posterior $p(\mathbf{G} \mid \mathcal{D})$ for linear Gaussian BNs, we follow the predominant choice in the literature and employ the Bayesian Gaussian Equivalent ( $B G e$ ) marginal likelihood, under which Markov equivalent structures are scored equally [16, 17]. Details on the computation of the BGe score are provided by Kuipers et al. [62]. Following the notation of Geiger and Heckerman [17] and Kuipers et al. [62], we use the standard effective sample size hyperparameters $\alpha_{\mu}=1$ and $\alpha_{\omega}=d+2$ as well as the diagonal form of the Wishart inverse scale matrix for the Normal-Wishart parameter prior underlying the BGe score.

Nonlinear The interaction between variables $\mathbf{x}$ can straightforwardly be extended to be nonlinear, e.g., using neural networks. In Section 6.3, we follow Zheng et al. [27] and consider (fully connected) feed-forward neural networks (FFNs) of the form

$$
\begin{align*}
& \operatorname{FFN}(\cdot ; \boldsymbol{\Theta}): \mathbb{R}^{d} \rightarrow \mathbb{R} \\
& \operatorname{FFN}(\mathbf{u} ; \boldsymbol{\Theta}):=\boldsymbol{\Theta}^{(L)} \sigma\left(\ldots \boldsymbol{\Theta}^{(2)} \sigma\left(\boldsymbol{\Theta}^{(1)} \mathbf{u}+\boldsymbol{\theta}_{b}^{(1)}\right)+\boldsymbol{\theta}_{b}^{(2)} \ldots\right)+\boldsymbol{\theta}_{b}^{(L)} \tag{E.2}
\end{align*}
$$

with weights $\boldsymbol{\Theta}^{(l)} \in \mathbb{R}^{d_{l} \times d_{l-1}}$, biases $\boldsymbol{\theta}_{b}^{(l)} \in \mathbb{R}^{d_{l}}$, and elementwise activation function $\sigma: \mathbb{R} \rightarrow \mathbb{R}$. Zheng et al. [27] show that the class of fully connected neural networks in (E.2) that do not depend on the value of $u_{k}$ is equivalent to the class of fully connected neural networks in (E.2) where the $k$-th column of $\Theta^{(1)}$ equals zero. This insight allows us to define a nonlinear Gaussian BN parameterized by a fully connected neural network:

$$
\begin{equation*}
p(\mathbf{x} \mid \mathbf{G}, \boldsymbol{\Theta})=\prod_{i=1}^{d} \mathcal{N}\left(x_{i} ; \operatorname{FFN}\left(\mathbf{G}_{i}^{\top} \circ \mathbf{x} ; \mathbf{\Theta}_{i}\right), \sigma^{2}\right) \tag{E.3}
\end{equation*}
$$

As required for a BN , each variable is independent of its non-descendants given its parents. The mask representation in (E.1) and (E.3) is equivalent to the concept of a structural gate used by Kalainathan et al. [63]. Note that the conditional distribution parameters for a single nonlinear Gaussian BN of the form in (E.3) contain the weights and biases of $d$ different neural networks, one for the local conditional distribution of each node.

## E. 2 Evaluation metrics

We provide additional details on the evaluation metrics used throughout the paper. Bayesian structure learning beyond five variables is notoriously difficult to evaluate since the ground truth posterior is not accessible. We hence rely and build on the metrics established in the literature.

Expected structural Hamming distance The structural Hamming distance $\operatorname{SHD}\left(\mathbf{G}, \mathbf{G}^{*}\right)$ between two graphs $\mathbf{G}$ and $\mathbf{G}^{*}$ counts the edge changes that separate the essential graphs representing the MECs of $\mathbf{G}$ and $\mathbf{G}^{*}[8,55]$. We define the expected structural Hamming distance to the ground truth graph $\mathbf{G}^{*}$ under the inferred posterior as

$$
\mathbb{E}-\operatorname{SHD}\left(p, \mathbf{G}^{*}\right):=\sum_{\mathbf{G}} p(\mathbf{G} \mid \mathcal{D}) \cdot \operatorname{SHD}\left(\mathbf{G}, \mathbf{G}^{*}\right) .
$$

Empirically, the $\mathbb{E}$-SHD is similar to the $L_{1}$ edge error used by Tong and Koller [11] and Murphy [10], but also takes into account the MEC. The $\mathbb{E}$-SHD is computed via Monte Carlo estimation of the expectation using samples from the posterior. Note that the DAG bootstrap variants and DiBS+ use the weighted mixture rather than the empirical distribution of samples. In the joint inference setting, we empirically marginalize out $\boldsymbol{\Theta}$ to obtain $p(\mathbf{G} \mid \mathcal{D})$.

Receiver operating characteristic The marginal posterior $p(\mathbf{G} \mid \mathcal{D})$ provides a confidence estimate $p\left(g_{i j}=1 \mid \mathcal{D}\right)$ for whether a given edge $(i, j)$ is present in the ground truth DAG G*. Recall that the marginal posterior edge probability $p\left(g_{i j}=1 \mid \mathcal{D}\right)$ is the posterior mean of an indicator for the presence of that edge, i.e., $p\left(g_{i j}=1 \mid \mathcal{D}\right)=\mathbb{E}_{p(\mathbf{G} \mid \mathcal{D})} \mathbb{1}\left[g_{i j}=1\right]$, which amounts to counting the proportion of graphs with $g_{i j}=1$ (and to weighted counting for the DAG bootstrap variants and DiBS+). The receiver operating characteristic (ROC) curve is then obtained by viewing the presence of each of the $d^{2}$ possible edges in a $d$-node graph as a binary classification task and varying the decision threshold from 0 to 1 under our confidence estimates $p\left(g_{i j}=1 \mid \mathcal{D}\right)$. The area under the receiver operating characteristic curve (AUROC) evaluates faithful uncertainty quantification of the posterior. In general, random guessing achieves an AUROC of 0.5 in expectation; a perfect classifier achieves an AUROC of 1.

Held-out log likelihood We also evaluate the ability to predict future observations by computing the average negative log likelihood on 100 held-out observations $\mathcal{D}^{\text {test }}$ defined as

$$
\text { neg. } \operatorname{LL}\left(p, \mathcal{D}^{\text {test }}\right):=-\sum_{\mathbf{G}, \boldsymbol{\Theta}} p(\mathbf{G}, \boldsymbol{\Theta} \mid \mathcal{D}) \cdot \log p\left(\mathcal{D}^{\text {test }} \mid \mathbf{G}, \boldsymbol{\Theta}\right)
$$

As for $\mathbb{E}$-SHD, the neg. LL is a posterior mean and thus computed via Monte Carlo estimation using samples from the posterior. When inferring $p(\mathbf{G} \mid \mathcal{D})$, the corresponding neg. MLL metric uses the marginal likelihood $p\left(\mathcal{D}^{\text {test }} \mid \mathbf{G}\right)$ instead of the likelihood $p\left(\mathcal{D}^{\text {test }} \mid \mathbf{G}, \boldsymbol{\Theta}\right)$.

Held-out log interventional likelihood Lastly, to capture relevant performance metrics in causal inference [10,12], we also compute the negative interventional log likelihood. Given an interventional
data set $\left(\mathcal{D}^{\text {int }}, \mathcal{I}\right)$, the interventional likelihood is given by

$$
\begin{equation*}
p\left(\mathcal{D}^{\mathrm{int}} \mid \mathbf{G}, \boldsymbol{\Theta}, \mathcal{I}\right)=\prod_{\mathbf{x}^{(n)} \in \mathcal{D}} \prod_{\substack{j=1 \\ j \notin \mathcal{I}}} p\left(x_{j}^{(n)} \mid \mathbf{x}_{\mathbf{G}_{j}}^{(n)}, \boldsymbol{\theta}_{j}\right) \tag{E.4}
\end{equation*}
$$

where $\mathbf{x}_{\mathbf{G}_{j}}$ are the values of the parents of variable $j$ in $\mathbf{G}$, and $\boldsymbol{\theta}_{j}$ parameterizes the local conditional distribution of $j$. The neg. I-LL and neg. I-MLL metrics are defined analogous to the neg. LL and neg. MLL in (19) but use the interventional likelihood in (E.4) instead of the observational likelihood. For marginal posterior inference, we likewise use the interventional marginal likelihood $p\left(\mathcal{D}^{\text {int }} \mid \mathbf{G}, \mathcal{I}\right)$ instead. In our experiments, we obtain interventional data $\left(\mathcal{D}^{\text {int }}, \mathcal{I}\right)$ by randomly selecting $10 \%$ of the variables and clamping them to zero in the ground-truth data-generating process. The reported neg. I-LL and I-MLL scores are the average of 10 different interventional data sets with $\left|\mathcal{D}^{\text {int }}\right|=100$.

## E. 3 Hyperparameters

In all evaluations, DiBS is run for 3,000 iterations and uses the simple linear constraint schedule $\beta_{t}:=t$. At $t=0$, the initial latent particles $\left\{\mathbf{Z}_{0}\right\}_{m=1}^{M}$ and parameter particles $\left\{\boldsymbol{\Theta}_{0}\right\}_{m=1}^{M}$ are initialized by sampling from their prior distributions. For the step size schedule $\eta_{t}$, we use the adaptive learning rate method RMSProp with learning rate 0.005 . We always use 128 samples for Monte Carlo estimation of the gradients. Finally, the bandwidths $\gamma_{z}, \gamma_{\theta}$ of the kernel in (15) and the slope of a linear schedule $\alpha_{t}$ are chosen in separate held-out instances of each setting in Section 6 and are listed in Table 2. As illustrated by the application in Section 7, the provided hyperparameters can be expected to apply to problem settings of comparable magnitude.

While the latent variable scale $\sigma_{z}$ can in principle be set arbitrarily, we always set $\sigma_{z}=1 / \sqrt{k}$ in the prior $p(\mathbf{Z})$, which makes the norm in the SE kernel given in (15) roughly invariant with respect to the latent dimension $k$, ignoring the acyclicity term. This follows from the fact that $\|\mathbf{u}\|^{2} \sim \operatorname{Gamma}\left(k / 2,2 \sigma_{z}^{2}\right)$ when $u_{i} \sim \mathcal{N}\left(0, \sigma_{z}^{2}\right)$, in which case $\mathbb{E}\left[\|\mathbf{u}\|^{2}\right]=k \sigma_{z}^{2}$.

Table 2: DiBS hyperparameter choices for $\alpha_{t}$ and bandwidths $\gamma_{z}, \gamma_{\theta}$. Here, $\widetilde{\alpha}$ denotes the slope in the linear schedule $\alpha_{t}:=\widetilde{\alpha} t$.

| Model | $d$ | $\widetilde{\alpha}$ | $\gamma_{z}$ | $\gamma_{\theta}$ |
| :--- | :---: | :---: | :---: | :---: |
| BGe | 20 | 2 | 2 | - |
|  | 50 | 2 | 50 | - |
| Linear Gaussian | 20 | 0.2 | 5 | 500 |
|  | 50 | 0.02 | 15 | 1,000 |
| Nonlinear Gaussian | 20 | 0.02 | 5 | 1,000 |
|  | 50 | 0.01 | 15 | 2,000 |

## E. 4 Baseline Methods

Structure MCMC ( $\mathbf{M C}^{3}, \mathbf{M}-\mathbf{M C}^{3}, \mathbf{G}-\mathbf{M C}^{3}$ ) Designed for inference of the marginal posterior $p(\mathbf{G} \mid \mathcal{D})$, structure MCMC $[36,37]$ performs sampling in the space of DAGs by adding, deleting, and reversing one edge at a time without violating acyclicity. The acceptance probability of a proposed graph $\mathbf{G}^{\prime}$ is given by

$$
\begin{equation*}
\min \left\{1, \frac{|\mathcal{N}(\mathbf{G})| \cdot p\left(\mathcal{D} \mid \mathbf{G}^{\prime}\right) p\left(\mathbf{G}^{\prime}\right)}{\left|\mathcal{N}\left(\mathbf{G}^{\prime}\right)\right| \cdot p(\mathcal{D} \mid \mathbf{G}) p(\mathbf{G})}\right\} \tag{E.5}
\end{equation*}
$$

where $\mathbf{G}$ is the current particle and $\mathcal{N}(\mathbf{G})$ is the collection of DAGs reachable from $\mathbf{G}$ with one edge change. Following [37], the ratio of neighborhoods is approximated to equal one, which allows for only computing $\mathcal{N}\left(\mathbf{G}^{\prime}\right)$ when accepting $\mathbf{G}^{\prime}$. We implement $\mathrm{MC}^{3}$ using the efficient ancestor matrix trick for finding acyclic proposals [37]. For marginal inference under the BGe marginal likelihood, we compute the Bayes factor in (E.5) by only taking into account the affected node families.

For all of $\mathrm{MC}^{3}, \mathrm{M}-\mathrm{MC}^{3}$, and G-MC ${ }^{3}$, we specify a burn-in period of 100 k samples and then collect a sample every 10k steps, which makes the wall time of $\mathrm{MC}^{3}$ and DiBS on CPUs comparable. Both $\mathrm{M}-\mathrm{MC}^{3}$ and G-MC ${ }^{3}$ use a simple Gaussian random walk proposal for the parameters, respectively, with scale selected to roughly obtain an acceptance rate of 0.2 in each setting [64], when feasible in combination with the graph proposal.

Nonparametric DAG bootstrap (BPC, BGES, BPC*, BGES*) The nonparametric DAG bootstrap [42] performs model averaging by bootstrapping the observations $\mathcal{D}$ to yield a collection of synthetic data sets, each of which is used to learn a single graph, here using the GES and PC algorithms [6, 7]. The collection of unique single graphs approximates the posterior by weighting each graph by its unnormalized posterior probability in (2), analogous to DiBS+. The closed-form maximum likelihood parameter estimate for linear Gaussian BNs with known G, which is used by BPC* and BGES* to allow approximating the joint posterior, is provided by Hauser and Bühlmann [57]. For joint posterior inference, $\mathrm{BPC}^{*}$ and $\mathrm{BGES}^{*}$ use $p(\mathbf{G}, \boldsymbol{\Theta}, \mathcal{D})$ rather than $p(\mathbf{G}, \mathcal{D})$ for weighting the inferred BN models.

Since the GES and PC algorithms only return essential graphs, i.e., MECs, we favor them in computing the AUROC score. We orient a predicted undirected edge correctly when a ground truth edge exists and only count a falsely predicted undirected edge as a single mistake. The held-out likelihood metrics given in (19) are computed for a random consistent DAG extension of the essential graph [65]. Enumerating the possibly exponential number of DAGs in an MEC is infeasible in general [66]. Implementations of the PC and GES algorithms are given by the CausalDiscoveryToolbox [67], which is published under an MIT Licence and executes their commonly used R implementations.

## F Efficient Implementation and Computing Resources

DiBS and SVGD operate on continuous tensors and solely rely on Monte Carlo estimation and parallel gradient ascent-like updates. Thus, our inference framework allows for a highly efficient implementation using vectorized operations, automatic differentiation, just-in-time compilation, and hardware acceleration. For this purpose, we implement DiBS with JAX [68], which is published under an Apache Licence. Our code is publicly available at: https://github.com/larslorch/dibs.

Table 3 summarizes the computing time of DiBS on GPU and CPUs for a superset of the inference tasks on BNs with Erdős-Rényi structures in Section 6. The GPU wall times for medium-sized inference tasks of up to around 20 nodes and 30 particles lie on the order of seconds or a few minutes. None of the baseline methods considered in this work are comparable in terms of efficiency and usage of modern hardware accelerators. Moreover, in larger inference problems than evaluated here, JAX would directly allow for the computations of DiBS to be performed in a distributed fashion, e.g., by updating batches of SVGD particles across multiple GPU devices. Note that BGe wall times are relatively slow because the closed-form marginal likelihood involves determinants [16, 17].

Table 3: Wall times of DiBS for the hyperparameters described in Section E.3. Times are the mean of 10 random restarts. $M$ denotes the number of particles, $d$ the number of nodes, and GPU/CPU the processing backend of JAX. We used one NVIDIA GEFORCE RTX 2080 Ti GPU or one full 2.70 GHz InTEL XEONGOLD 6150 CPU node to measure GPU and CPU wall time, respectively, for each run. Experiments marked by a dash exceeded the GPU memory. The main experiments of Sections 6 and 7 were performed in bulk on Oracle BM.STANDARD.E2.64 CPU machines and no GPUs.

| Model $\quad d=$ | Wall time (min) |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | GPU |  |  | CPU |  |  |
|  | 10 | 20 | 50 | 10 | 20 | 50 |
| BGe |  |  |  |  |  |  |
| $M=10$ | 0.349 | 0.892 | 8.111 | 3.337 | 18.251 | 216.844 |
| $M=30$ | 0.609 | 2.084 | - | 9.380 | 52.203 | 659.106 |
| Linear Gaussian |  |  |  |  |  |  |
| $M=10$ | 0.370 | 0.603 | 1.571 | 1.410 | 3.738 | 21.531 |
| $M=30$ | 0.612 | 1.153 | 3.975 | 4.703 | 13.163 | 74.914 |
| Nonlinear Gaussian |  |  |  |  |  |  |
| $M=10$ | 2.702 | 6.139 | 24.117 | 10.019 | 27.749 | 128.820 |
| $M=30$ | 7.476 | 17.667 | - | 28.130 | 79.628 | 388.992 |

## G Results for Gaussian Bayesian networks with $d=50$ variables





Figure 5: Marginal posterior inference of linear Gaussian BNs with $d=50$ variables using the BGe marginal likelihood. The metrics are aggregated for 30 random BNs of each graph type. While DiBS and DiBS+ are competitive in the structural $\mathbb{E}$-SHD and AUROC metrics, we find that the baselines specifically designed for marginal posterior inference perform favorably in the likelihood-based metrics. We hypothesize that this is due to the high variance incurred by the score function estimator that DiBS needs to use in the marginal inference setting under the BGe model (cf. Section 6.2). To reach comparable results with DiBS in this high-dimensional setting, the DiBS score function gradient estimator may require more than the default 128 Monte Carlo samples used here.


Figure 6: Joint posterior inference of linear Gaussian BNs with $d=50$ variables. The first and second rows show the aggregate metrics for inference of 30 random BNs with Erdős-Rényi and scale-free structures, respectively. Analogous to inference for $d=20$ variables, DiBS+ outperforms all alternatives to joint posterior inference of the graph and the conditional distribution parameters across the metrics.


Figure 7: Joint posterior inference of nonlinear Gaussian BNs with $d=50$ variables, where each local conditional distribution is parameterized by a 2-layer neural network with five hidden nodes. In this setting, the total number of conditional distribution parameters in a given BN amounts to $|\boldsymbol{\Theta}|=13,050$ weights and biases. The metrics are aggregated for inference of 30 random BNs of each graph type. Here, DiBS only infers 10 particles to make the wall time on CPUs comparable to M-MC ${ }^{3}$ and G-MC ${ }^{3}$. As for posterior inference of BNs with $d=20$ variables, DiBS and DiBS+ perform favorably compared to the $\mathrm{MC}^{3}$ baselines.

## H Additional Analyses and Ablation Studies

Having compared DiBS with several alternative approaches to Bayesian structure learning in Section 6 , this supplementary section is devoted to a more in depth analysis of some of its properties. This is done by changing, or leaving out single design aspects of the algorithm and studying the effect on the previous metrics.

As in Section 6, DiBS and its instantiation with SVGD are used interchangeably here, and DiBS+ denotes the weighted mixture of particles. Since the metrics do not qualitatively differ between inference of Erdős-Rényi and scale-free BN structures in our experiments of Section 6, we only consider the former here. Unless mentioned otherwise, the following experimental setup corresponds to joint posterior inference of linear Gaussian BNs with $d=20$ variables in Section 6.2.

## H. 1 Graph Embedding Representation

In Section 4.2, we propose to use a generative graph model $p_{\alpha}(\mathbf{G} \mid \mathbf{Z})$ that is based on the inner product of latent embeddings for each node. In particular, we choose $p_{\alpha}\left(g_{i j}=1 \mid \mathbf{Z}\right)=\sigma_{\alpha}\left(\mathbf{u}_{i}^{\top} \mathbf{v}_{j}\right)$ with latent variables $\mathbf{Z}=[\mathbf{U}, \mathbf{V}]$. In Figure 8, we contrast this modeling choice with the more trivial variant $p_{\alpha}\left(g_{i j}=1 \mid \mathbf{Z}\right)=\sigma_{\alpha}\left(\mathbf{z}_{i j}\right)$, where single scalars rather than inner products between latent


Figure 8: Contrasting the bilinear graph model of Section 4.2 with its more trivial variant, where each latent variable models the edge probabilities directly via the sigmoid. The plots aggregate the results for joint inference of 30 randomly generated linear Gaussian BNs with $d=20$ variables.


Figure 9: DiBS for joint inference of linear Gaussian BNs with $d=20$ variables for different sizes of the latent variables $\mathbf{Z} \in \mathbb{R}^{2 \times d \times k}$. Lower rank parameterizations of the matrix of edge probabilities balance the tradeoff between computational efficiency and posterior approximation quality.
vectors encode the edge probabilities. Bengio et al. [69] and Ke et al. [31] use the scalar variant with fixed $\alpha=1$ in the context of causal inference.
The comparison in Figure 8 illustrates that incorporating only the bilinear parameterization of edge probabilities in the generative graph model improves performance by a significant margin. We hypothesize that the coupling between edges results in smoother densities, which might be less prone to local minima in gradient-based methods such as DiBS.

## H. 2 Graph Embedding Dimensionality

Another feature of the inner product representation of graphs is the ability to control the dimensionality of the posterior inference task. As described in Section 5, we generally set $k=d$ for the latent variables $\mathbf{Z} \in \mathbb{R}^{2 \times d \times k}$ that parameterize our graph model $p_{\alpha}(\mathbf{G} \mid \mathbf{Z})$. This leaves the matrix of edge probabilities fully expressible and without a rank constraint. In principle, however, the formulation in (6) allows us to arbitrarily vary $k$. This creates a trade-off between the complexity of the parameterization and the tractability and dimensionality challenges in approximate inference of $p(\mathbf{Z} \mid \mathcal{D})$ or $p(\mathbf{Z}, \boldsymbol{\Theta} \mid \mathcal{D})$. Limiting $k<d$ has connections to the theory of low-rank realizations of sign matrices.
We perform inference with $\operatorname{DiBS}$ for $k \in\{5,7,10,15,20\}$, leaving all other aspects of the algorithm unchanged. Hence, the corresponding posterior over $\mathbf{Z}$ has $\{200,280,400,600,800\}$ dimensions, respectively. The results in Figure 9 suggest that lower values of $k=15$, or even $k=10$, are already able to achieve competitive performance across all metrics. Interestingly, the structural $\mathbb{E}$-SHD metric appears to suffer most from a small loss in complexity.
In this context, one should keep in mind that the bandwidth parameters $\gamma_{z}$ and $\gamma_{\theta}$ were set to achieve good performance with $k=d=20$. It is possible that lower values of $k$ can reach performances that are even closer to the full-rank variant of DiBS with alternative settings for $\gamma_{z}$ and $\gamma_{\theta}$. In addition, a lower-rank DiBS variant could be particularly promising for inference of very large BNs, where the


Figure 10: Performance of DiBS and $\mathrm{DiBS}+$ as a function the number of particle transport steps $T$. As previously, the plots aggregate the results for inference of 30 randomly generated linear 20-node Gaussian BNs. The latent variable $\mathbf{Z}$ is specified with its default dimensions $k=d=20$. After already roughly 1,000 iterations, DiBS and $\mathrm{DiBS}+$ obtain good posterior approximations.
computational challenges of the full $O\left(d^{2}\right)$ latent representation might outweigh its benefits in terms of expressibility.

## H. 3 Particle Transport Iterations

Since DiBS uses Stein variational gradient descent [19] for posterior inference, our method iteratively transports a set of latent graph particles, or latent graph and parameter tuples, for a number of $T$ steps. As the particles are randomly initialized, the approximation quality of SVGD particles (and thus also DiBS particles) improves with the number of steps. We are interested in the degree to which a small number of transport steps provide a good posterior approximation in the face of computational constraints.

In Figure 10, we show the performance of DiBS as a function of the number of transport steps. We find that even a smaller number of iterations achieves competitive results across the metrics. In addition, the variance of performance in the predictive metrics neg. LL and neg. I-LL decreases monotonically as a function of the performed transforms, whereas the variation in $\mathbb{E}$-SHD remains roughly the same.

## H. 4 Uncertainty Quantification Within a Markov Equivalence Class

When performing posterior inference of $p(\mathbf{G} \mid \mathcal{D})$ with the BGe marginal likelihood and a uniform prior $p(\mathbf{G})$, each Markov equivalent structure is assigned equal likelihood. This might be desirable considering that causal edge directions are often not fully identifiable from purely observational data. As DiBS infers a posterior over DAGs rather than MECs, we aim to validate the ability of DiBS to correctly quantify the uncertainty present in nonidentifiable edge directions.

To this end, we consider a 4-node example Bayesian network, small enough to allow for the closedform computation of the ground-truth posterior by exhaustive enumeration of all possible DAGs. This enables us to compute the true single and pairwise posterior edge marginals and contrast them with the approximate posterior marginals inferred by DiBS. The graph structure for this analysis is chosen to contain both an identifiable v-structure and a nonidentifiable edge pair. Figure 11 shows the ground truth DAG GG${ }_{0}^{*}$, its linear Gaussian parameters, and the observational model. In addition,


$$
\begin{align*}
p\left(x_{0}\right) & =\mathcal{N}\left(x_{0} ; 0,1\right) \\
p\left(x_{1} \mid x_{0}\right) & =\mathcal{N}\left(x_{1} ; 2 x_{0}, 1\right) \\
p\left(x_{2} \mid x_{0}\right) & =\mathcal{N}\left(x_{2} ;-2 x_{0}, 1\right)  \tag{H.1}\\
p\left(x_{3} \mid x_{1}, x_{2}\right) & =\mathcal{N}\left(x_{3} ; 3 x_{1}+x_{2}, 1\right)
\end{align*}
$$



Figure 11: Four-node example linear Gaussian Bayesian network. Under the BGe marginal likelihood and a uniform prior, $\mathbf{G}_{0}^{*}, \mathbf{G}_{1}^{*}$, and $\mathbf{G}_{2}^{*}$ are scored equally. While the $v$-structure $x_{1} \rightarrow x_{3} \leftarrow x_{2}$ is an identifiable feature of the MEC of $\mathbf{G}_{0}^{*}$ and thus present in $\mathbf{G}_{\mathrm{MEC}}^{*}$, the edge directions of $x_{1}-x_{0}-x_{2}$ cannot be distinguished even given infinite observational data.

Table 4: Ground truth and average inferred posterior marginals given $N=100$ observations from the ground truth model in Figure 11. Listed are the probabilities for the nonidentifiable edge structure $x_{1}-x_{0}-x_{2}$ (top) and the identifiable v-structure $x_{1} \rightarrow x_{3} \leftarrow x_{2}$ (bottom). Averaged over 30 random particle initilizations, DiBS+ correctly quantifies the confidence and uncertainty in the v -structure and nonidentifable edge pair, respectively.

|  | DiBS | DiBS+ | Ground Truth |
| :--- | :---: | :---: | :---: |
| $p\left(x_{1} \rightarrow x_{0}, x_{0} \rightarrow x_{2} \mid \mathcal{D}\right)$ | 0.134 | 0.216 | $\mathbf{0 . 2 9 8}$ |
| $p\left(x_{1} \rightarrow x_{0}, x_{0} \leftarrow x_{2} \mid \mathcal{D}\right)$ | 0.201 | 0.037 | $\mathbf{0 . 0 5 2}$ |
| $p\left(x_{1} \leftarrow x_{0}, x_{0} \rightarrow x_{2} \mid \mathcal{D}\right)$ | 0.223 | 0.409 | $\mathbf{0 . 3 1 1}$ |
| $p\left(x_{1} \leftarrow x_{0}, x_{0} \leftarrow x_{2} \mid \mathcal{D}\right)$ | 0.103 | 0.298 | $\mathbf{0 . 2 9 7}$ |
| $p\left(x_{1} \rightarrow x_{3}, x_{3} \rightarrow x_{2} \mid \mathcal{D}\right)$ | 0.157 | 0.013 | $\mathbf{0 . 0 1 7}$ |
| $p\left(x_{1} \rightarrow x_{3}, x_{3} \leftarrow x_{2} \mid \mathcal{D}\right)$ | 0.338 | 0.934 | $\mathbf{0 . 9 1 4}$ |
| $p\left(x_{1} \leftarrow x_{3}, x_{3} \rightarrow x_{2} \mid \mathcal{D}\right)$ | 0.275 | 0.034 | $\mathbf{0 . 0 4 3}$ |
| $p\left(x_{1} \leftarrow x_{3}, x_{3} \leftarrow x_{2} \mid \mathcal{D}\right)$ | 0.154 | 0.019 | $\mathbf{0 . 0 2 5}$ |

Figure 11 lists the essential graph $\mathbf{G}_{\mathrm{MEC}}^{*}$ as well as the two other Markov equivalent DAGs $\mathbf{G}_{1}^{*}$ and $\mathrm{G}_{2}^{*}$ in the MEC represented by $\mathrm{G}_{\mathrm{MEC}}^{*}$.
We perform marginal posterior inference with DiBS using the experimental setup and hyperparameters for 20-node linear Gaussian BNs. In this example setting, DiBS employs a uniform prior over graphs and uses the default $k=d=4$. Table 4 shows the ground truth and inferred pairwise edge marginals under the posterior. We find that $\mathrm{DiBS}+$ correctly infers both the uncertainty in the edge directions of $x_{1}-x_{0}-x_{2}$ as well as the high confidence in the presence of the $v$-structure $x_{1} \rightarrow x_{3} \leftarrow x_{2}$. While the unweighted particles of DiBS do not exhibit false confidence in structures that are not present in $\mathbf{G}_{0}^{*}$, its inferred degree of uncertainty is too high compared to the ground truth. The DiBS+ variant overcomes the inexact empirical average of DiBS by weighting the particles by their unnormalized posterior probabilities.

## I Experimental Details for Application to Protein Signaling Networks

The data by Sachs et al. [3] as well as the corresponding consensus graph used in Section 7 are taken as provided by the CausalDiscoveryToolbox [67], which is published under an MIT Licence. We standardize the data for inference. Because $N$ is large, DiBS uses minibatches of 100 observations to estimate the scores of the posterior. All hyperparameters and BN specifications are chosen by default exactly as during the synthetic evaluation of linear and nonlinear Gaussian BNs in Table 2, respectively, except that DiBS correspondingly uses $k=d=11$. For joint posterior inference of linear Gaussian BNs, BPC* and BGES* still use the BGe marginal likelihood; since metrics are very similar to BPC and BGES, their scores are not reported.
In line with inference on synthetic data in Section 6, the BGe marginal likelihood employed for the experiments in Section 7 uses the default effective sample sizes $\alpha_{\mu}=1$ and $\alpha_{\omega}=d+2$ described in Appendix E.1. Likewise, we again set the noise level for inference with the explicitly parameterized linear and nonlinear Gaussian networks to $\sigma^{2}=0.1$.

Since the effective sample size $\alpha_{\mu}$ and the noise level $\sigma^{2}$ may affect the model complexity of the inferred BNs, e.g., the mean number of inferred edges in the DAG, we provide additional results for alternative values of these Bayesian network model hyperparameters in Tables 5 and 6. Overall, we find that increasing the effective sample size $\alpha_{\mu}$ does not significantly change metrics across the considered methods. However, higher fixed noise levels $\sigma^{2}$ do result in less inferred edges, which tends to lead to lower $\mathbb{E}$-SHD but worse AUROC, i.e., less calibrated edge confidence scores. We note that these are not free parameters of the inference methods that approximate the posterior, but specifications of the inferred BN models themselves.

Table 5: Additional results for marginal posterior inference of protein signaling pathways under the BGe marginal likelihood of linear Gaussian BNs [16, 17]. Changing the effective sample size in the BGe Normal-Wishart prior does not result in significantly different metrics compared to the default $\alpha_{\mu}=1$ used in all of our experiments. For $\alpha_{\mu}=10$, DiBS and DiBS+ average an expected number of 39.6 and 35.4 edges, respectively. Metrics are the mean $\pm \mathrm{SD}$ of 30 random restarts.

|  | $\alpha_{\mu}=10$ |  |
| :--- | :---: | :---: |
|  | $\mathbb{E}-\mathbf{S H D}$ | AUROC |
| MC $^{3}$ | $34.3 \pm 0.4$ | $0.622 \pm 0.020$ |
| BPC | $25.5 \pm 2.3$ | $0.566 \pm 0.020$ |
| BGES | $33.8 \pm 1.8$ | $0.641 \pm 0.034$ |
| DiBS | $37.9 \pm 0.5$ | $0.637 \pm 0.046$ |
| DiBS + | $35.1 \pm 1.8$ | $0.627 \pm 0.050$ |

Table 6: Additional results for joint posterior inference of protein signaling pathways under explicitly parameterized linear (top) and nonlinear (bottom) Gaussian BNs. The hyperparameter $\sigma^{2}$ specifies the noise level underlying the inferred Bayesian networks. We find that higher noise levels $\sigma^{2}$ tend to result in less edges. When inferring linear Gaussian BNs with $\sigma^{2}=0.01\left(\sigma^{2}=1\right)$, DiBS averages an expected number of 11.6 (8.8) edges, DiBS+ 13.8 (9.9) edges. For nonlinear Gaussian BNs with $\sigma^{2}=0.01\left(\sigma^{2}=1\right)$, $\operatorname{DiBS}$ averages $15.6(5.2)$ edges, $\operatorname{DiBS}+17.5(6.8)$ edges. Due to less false positives, the $\mathbb{E}$-SHD improves, but the degree of uncertainty in the presence of edges is quantified less accurately, resulting in worse AUROC. Metrics are the mean $\pm \mathrm{SD}$ of 30 random restarts.

|  | $\sigma^{2}=0.01$ |  | $\sigma^{2}=1$ |  |
| :--- | :---: | :---: | :---: | :---: |
|  | E-SHD | AUROC | $\mathbb{E}-$ SHD | AUROC |
| M-MC $^{3}$ | $38.1 \pm 3.4$ | $0.536 \pm 0.082$ | $33.1 \pm 3.4$ | $0.543 \pm 0.105$ |
| G-MC | 3 | $30.9 \pm 3.0$ | $0.518 \pm 0.051$ | $29.8 \pm 3.7$ |
| DiBS | $23.0 \pm 0.5$ | $0.531 \pm 0.078$ |  |  |
| DiBS + | $22.9 \pm 2.0$ | $0.540 \pm 0.069$ | $20.3 \pm 0.4$ | $0.601 \pm 0.039$ |
|  | $\sigma^{2}=0.01$ |  | $20.0 \pm 1.4$ | $0.569 \pm 0.040$ |
|  | $\mathbb{E}-$ SHD | AUROC | $\sigma^{2}=1$ |  |
| M-MC $^{3}$ | $38.7 \pm 3.2$ | $0.555 \pm 0.101$ | $18.4 \pm 0.1$ | $0.501 \pm 0.043$ |
| G-MC $^{3}$ | $34.9 \pm 3.6$ | $0.542 \pm 0.064$ | $30.6 \pm 2.5$ | $0.538 \pm 0.059$ |
| DiBS | $24.3 \pm 0.6$ | $0.582 \pm 0.050$ | $17.7 \pm 0.1$ | $0.550 \pm 0.020$ |
| DiBS+ | $24.9 \pm 2.9$ | $0.535 \pm 0.045$ | $18.5 \pm 0.5$ | $0.530 \pm 0.028$ |


[^0]:    ${ }^{1}$ Our Python JAX implementation of DiBS is available at: https://github.com/larslorch/dibs

